DEPARTMENT OF ELECTRICAL & COMPUTER ENGINEERING
COLLEGE OF ENGINEERING & TECHNOLOGY
OLD DOMINION UNIVERSITY
NORFOLK, VIRGINIA 23529

**SIGNAL PROCESSING METHODOLOGIES FOR AN
ACOUSTIC FETAL HEART RATE MONITOR**

P. 105

By

Robert A. Pretlow III, Graduate Research Assistant

and

John W. Stoughton, Principal Investigator

Final Report
For the period January 1, 1988 to September 30, 1991

September 1992

DEPARTMENT OF ELECTRICAL & COMPUTER ENGINEERING
COLLEGE OF ENGINEERING & TECHNOLOGY
OLD DOMINION UNIVERSITY
NORFOLK, VIRGINIA 23529

SIGNAL PROCESSING METHODOLOGIES FOR AN
ACOUSTIC FETAL HEART RATE MONITOR

By

Robert A. Pretlow III, Graduate Research Assistant

and

John W. Stoughton, Principal Investigator

Final Report
For the period January 1, 1988 to September 30, 1991

September 1992

EXECUTIVE OVERVIEW

This document is the final report for research conducted for Cooperative Agreement NCC1-120 which ended July 31, 1991. In this report, the design and preliminary evaluation of a passive acoustic fetal heart rate monitor is described. The detailed description of this research effort follows this section and is taken from the masters' thesis titled "Signal Processing Methodologies for an Acoustic Fetal Heart Rate Monitor" by Robert A. Pretlow, III, M.D.. Dr. Pretlow's research was directly supported by this cooperative agreement.

The system development was an effort to incorporate NASA technology into a fetal heart rate monitor as first conceptualized by Dr. Donald Baker. Of interest is the use of polyvinyl fluoride (PDVF) piezoelectric film which is used to monitor vibrations in wind tunnel models, as sensitive microphone for the pick up of fetal heart tones. The cooperative agreement between ODU and NASA was to develop a signal processing strategy as part of a passive acoustic fetal heart rate monitor. The purpose of the system is provide noninvasive fetal monitoring for high risk pregnancies with out the need for costly and perhaps unobtainable clinical evaluation. At present, the accepted approach to fetal monitoring is with ultrasound-based monitoring which may be administered only by specialists within a clinical setting. The form of monitoring being considered is the fetal nonstress test which determines the fetal well being by monitoring heart rate acceleration and deceleration after perceived fetal movement, such as when the fetus turns or kicks.

The prototype fetal heart rate system is characterized by a sensor belt consisting of an array of PDVF sensors which provide the acoustic pickup of the fetal heart tones. The acoustic signals are then transformed by the PVF2 material to electrical signals which then are then digitally sampled and operated upon by the digital signal processing unit. The digital signal processing unit is composed of a personal computer with an imbedded digital signal processing unit. In the signal processing unit, the fetal heart tone is detected using a linear prediction paradigm. The product of the signal processing is an estimate of the instantaneous heart rate which is then fed as an input to a strip chart recorder.

Preliminary clinical evaluation of the prototype system was performed with the cooperation of the Department of Fetal Maternal Medicine at Eastern Virginia Medical School (EVMS). Volunteers were solicited and gave prior consent under the

auspices of human subject regulations of both Old Dominion University and MCHR. The selected volunteers were in their thirty fifth to thirty ninth week of pregnancy. Each volunteer was given a normal ultrasound stress test followed by the recording of data with the prototype acoustic monitor system. Initial results showed that the prototype system fetal heart rate compared favorably to that recorded by the ultrasound system. The prototype acoustic monitor system has subsequently been turned over to NASA for evaluation by Dr. Donald Baker.

ABSTRACT

SIGNAL PROCESSING METHODOLOGIES FOR AN ACOUSTIC
FETAL HEART RATE MONITOR

Robert A. Pretlow, III
Old Dominion University, 1991
Director: Dr. John Stoughton

Research and development is presented of real time signal processing
methodologies for the detection of fetal heart tones within a noise-contaminated signal
from a passive acoustic sensor. A linear predictor algorithm is utilized for detection of
the heart tone event and additional processing derives heart rate. The linear predictor is
adaptively "trained" in a least mean square error sense on generic fetal heart tones
recorded from patients. A real time monitor system is described which outputs to a strip
chart recorder for plotting the time history of the fetal heart rate. The system is validated
in the context of the fetal nonstress test. Comparisons are made with ultrasonic nonstress
tests on a series of patients. Comparative data provides favorable indications of the
feasibility of the acoustic monitor for clinical use.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER ONE

INTRODUCTION

1.1 FETAL HEART RATE MONITORING

The fetal heartbeat was first detected in 1818 by a physician listening to a mother's abdomen with a crude funnel-shaped tube [1]. In 1833 a textbook on "Obstectric Ascultation" noted the possible relationship between certain fetal heart rate (FHR) patterns and fetal well-being [2]. Since that time the clinical utility of FHR monitoring has become well established as a means of assessing the health of the fetus [3]. FHR monitoring is performed during pregnancy in the form of the so-called nonstress test (NST). In the NST, if a 15-second long fetal heart rate acceleration of 15 beats per minute over its average baseline value follows a fetal movement three times during a 20 minute period, then the fetus is judged to be healthy [4]. Figure 1.1 illustrates the NST.

The NST is conventionally performed in a clinic or physician's office because of the size and complexity of the equipment. Monitoring of the fetus during pregnancy is thus possible for only brief, intermittent periods. Three methods have been used to monitor fetal heart rate and perform the NST. They are fetal electrocardiogram (FECG) techniques, continuous wave ultrasonic Doppler-shift techniques, and fetal phonocardiogram (FPCG) techniques.

1.1.1 FECG

FECG techniques detect the changing electrical field of the beating fetal heart and involve the attachment of electrodes either to the mother's abdomen or directly to the

1

fetus if rupture of the fetal membranes has occurred during labor. The maternal abdominal signal is of very low amplitude and highly contaminated with large amplitude maternal ECG and background noise and has not been found to be practical for FHR monitoring [40]. The method of attaching electrodes directly to the fetus is available only when the fetal membranes are ruptured and the mother is commited to delivery.

1.1.2 Doppler Ultrasonography

Doppler ultrasonography involves the placement of a continuous-wave ultrasound transducer on the maternal abdomen and detection of the Doppler frequency shift in the reflected signal produced by the moving structures of the fetal heart. It is a quite sensitive technique and allows detection of the fetal heart beat when membranes are intact. It is, however, theoretically an "invasive" technique in that the fetus is continually exposed to the energy of the ultrasonic beam, although no evidence of fetal harm has been demonstrated to date. Also, disruption of detection will occur if the fetal heart moves (with fetal movement) out of the path of the relatively narrow beam.

1.1.3 FPCG

FPCG techniques detect the fetal heartbeat sounds or fetal heart "tones" (FHT) by means of a passive microphone applied to the maternal abdomen. The FPCG technique was used in early fetal monitors but eventually was abandoned in favor of ultrasonography because of the greater sensitivity of the latter [5]. The main advantages of the acoustic technique are its passivity (non-invasiveness) and its simplicity.

1.2 Background of the NASA/ODU/EVMS Project

In 1986 Dr. Donald Baker of Spokane, Washington suggested that long term fetal monitoring of the mother at home might prevent a significant number of fetal problems such as hypoxic brain damage. He further suggested that a passive acoustic device

utilizing the FPCG signal would obviate the invasive risks inherent with long term ultrasonic monitoring. Moreover, the imagined simplicity of an acoustic fetal monitor would hopefully allow home use.

After attempting some development of this concept on his own Dr. Baker eventually made contact with the Technology Utilization Division of NASA whose charter is to make NASA's technology available to the private sector. NASA agreed to fund a three year project to develop Dr. Bakers's idea into a working monitor. NASA would fabricate the sensors for the monitor utilizating a polyvinyl fluoride acoustic film used to detect vibrations on windtunnel models. The Department of Electrical and Computer Engineering at Old Dominion University was enlisted to develop the signal processing hardware and software. The Department of Fetal Maternal Medicine at Eastern Virginia Medical School (EVMS) agreed to facilitate the clinical data collection and testing of the acoustic monitor.

1.3 Thesis Research Objective

The objective of this thesis research was to develop the real time signal processing methodologies which would 1) detect the FHT within the noise contaminated acoustic signal and 2) derive FHR. The system hardware and software subsequently would be tested with the sensor belt in the clinical context of the NST. The acoustic monitor FHR output would be compared with that of a conventional ultrasonic monitor on a series of patients.

1.4 Thesis Organization

The background and theory used to develop the heart tone detection algorithm is presented in Chapter Two. Chapter Two also includes the results of an initial data collection pilot study. The methodologies required to implement the detection algorithm in real time will be presented in Chapter Three along with the heart rate derivation

techniques. Also included in Chapter Three is a description of the system hardware. In Chapter Four the experimental NST study to validate the system hardware and software is presented. The FHR records of a series of patients undergoing simultaneous acoustic and ultrasonic NST's are compared. Chapter Five contains the experimental conclusions and a discussion of the developmental potential of the algorithm and the acoustic monitor.

CHAPTER TWO

BACKGROUND AND THEORY

2.1 Introduction

The background and theory for the design of the signal processing algorithm for the detection of fetal heart tones is presented in this chapter. In section 2.2 the characteristics of the FPCG signal are described, first as reported in the available literature and second as observed in the data collection pilot study of this research. Section 2.3 contains a review of previous heart beat event detection methodologies from the available literature. The suitability of these metodologies for the current research is discussed. Lastly, in section 2.4 the adaptively "trained" least mean square error linear predictor is described. The rationale for its use in the current application is discussed.

2.2 FPCG Signal Characteristics

The FPCG signal is a relatively low energy signal. Ascultation of the fetal heart tones by a physician generally requires a special type of stethescope, such as the deLee Hillis model which includes a metal bracket worn over the physician's head to increase sensitivity. Generally, the fetal heartbeat can be heard in only a small area of the mother's abdomen of usually no more than three centimeters (cm) radius, although the range of this local area can encompass up to a 12 cm radius as illustrated in figure 2.1.

The heart tone signal is thought to result from vibrations produced by the opening and closing of the four valves controlling blood flow through the fetal heart and from vibrations of the heart muscle. The heart tones are classified into two components, the

6

12 cm
Range

3 cm
Localization

Figure 2.1. Fetal Heart Tone Localization.

Figure 2.3. FPCG Signal (from Talbert et al.[6])

frequencies, stating, "the frequency spectrum ... has not been exactly measured because of dependence on the coupling condition between the transducer and abdominal wall." They observed the first heart sound at a frequency of 50 Hz and the second sound at 100 Hz. They claim to have determined the "true frequency spectrum of the fetal heart sounds" since their results include compensation for the response of their transducer on the abdominal wall. But their findings are suspect in that they passed the transducer signal through a 68 - 72 Hz bandpass filter prior to spectral analysis and their transducer was also tuned to resonate at 70 Hz. Additionally, they note that the frequency spectrum varies from one patient to another and also varies in the same fetus according to gestational age.

Nagel [11] used a transducer consisting of piezoelectric material bonded to a circular metal disk to determine the spectrum of the fetal heart tones. His results are shown in figure 2.4a, which depicts the power spectral density of the acoustic maternal abdominal signal. On a large series of patients he found the first fetal heart sound to be in the range of 20 Hz and the second sound about 45 Hz. Nagel also reports that the spectrum of the fetal heart sounds varies considerably with fetal gestational age as shown in figure 2.4b. He suggests the possibility of using the heart tone spectrum as a means of estimating fetal maturity.

The above literature findings are summarized in table 2.1.

2.2.2 FECG Signal

The fetal electrocardiogram (FECG) signal is produced by the electrical depolarization of the fetal heart corresponding to heart muscle contraction. The FECG signal is important because the FECG waveforms have many similarities to the FPCG waveforms. Also, considerable signal processing effort has been applied to processing the FECG, whereas there has been little effort applied to processing the FPCG.

(a)

(b)

Figure 2.4. Spectral Characteristics of FPCG (from Nagel [11]) :

a) Spectrum of Maternal Abdominal Signal.   b) FPCG Spectrum vs.
Gestational Age.

Table 2.1. Fetal Heart Frequencies per Available Literature

Source	First sound	Second sound
Hewlett-Packard [9]	80-110 Hz ⟶	
Jenssen [11]	40 Hz	40-50 Hz
Talbert et al. [12]	30 Hz	75-100 Hz
Talbert et al. [9]	60-80 Hz(40 Hz)	60 Hz
Kobayshi et al. [13]	50 Hz	100 Hz
Nagel [52]	20 Hz(45 Hz)	40-50 Hz

The FECG waveform consists of a small amplitude one half cycle "P" wave, from depolarization of the upper chambers (atria), followed by a large one and one half cycle "QRS" wave, from depolarization of the lower chambers (ventricles), followed by a small half cycle ventricular repolarization "T" wave. A typical FECG waveform and its time relationship to the FPCG waveform can be seen in figure 2.5.

2.2.3 Modeling of the FPCG Signal

The FPCG can be modeled as a sum of "almost" periodically recurring deterministic transients or "wavelets" [12],[13]. The interval between the wavelets is a random process, although its statistics have a certain degree of stationarity in healthy patients. The shape of the wavelet is fairly deterministic for a particular patient, although the wavelet shape and amplitude can vary slightly from one cycle to the next due to changes in heart blood flow dynamics.

The FPCG signal can be summarized by the following expression:

$$x(t) = s(t) + n(t)$$

where $x(t)$ is the composite signal consisting of the heart signal superimposed on background noise, $n(t)$ [13]. Background noise for the FPCG consists of maternal respiration, gastrointestinal, and muscle movement sounds; maternal heartbeat, aortic/placental pulse waves, and placental blood turbulence sounds; and environmental ambient noise and 60 Hz electromagnetic interference(EMI). The noise function is shown above as additive, but noise also can be multiplicative [14] as from variation of sensor coupling to the abdomen because of maternal movement and/or respiration, so that

$$x(t) = m(t)[s(t) + n(t)]$$

Figure 2.5.   Time Relationship of FPCG Signal to FECG Signal
(from Jenssen [8]).

where m(t) is the multiplicative noise function.

The maternal heart tones are of quite large amplitude as compared to the fetal acoustic signal but fortunately are of lower frequency range. There is some overlap of the spectrums, however. Nagel [11] found the maternal tones to be in the 8 - 15 Hz range as shown in figure 2-3a. The maternal aortic pulse wave is also a large amplitude signal to which the PDVF sensors of this research would be sensitive since their response is essentially down to DC . The spectral energy of this wave has not been reported in the literature, however, that of the general human pulse has been reported to have 99% of its spectral energy below 10Hz [15].

2.2.4 Data Collection Pilot Study

In order to formulate a heart tone detection strategy, a definitive knowledge of the fetal heart tone signature is essential as well as of the noise components of the raw sensor signal. Since the available literature is contradictory in this area, an initial data collection pilot study was conducted with the purpose of specifically identifying the fetal heart tone signature characteristics and the sources of noise contamination in the sensor signal. Additionally, the recorded data would be available for subsequent development and testing of the detection algorithm.

Fetal heart tone recordings were conducted on ten patients coming to the Department of Fetal-Maternal Medicine at EVMS for high risk obstectrical care but with no known fetal heart abnormalities. Approval for the human subjects study was obtained from the Institutional Review Boards of both ODU and EVMS. The recording system consisted of a triple sensor ("second generation") sensor belt (figure 4.2), a four channel medical isolation amplifier with a fixed gain of 100, two adjustable 4th order Butterworth high pass filters in series , a 6th order Bessel 167Hz anti-aliasing low pass filter, a second amplifier with a gain of 10 - 100 - 1000, a data aquisition board equipped with a sample/hold module, and a PC. The PC allowed wide-bandwidth recording of the

data on hard disk as well as display of the incoming signal waveforms on the monitor screen in real time. Additionally an audio amplifier with a set of headphones was used to optimally position the sensor belt.

2.2.5 Results of the Data Collection Pilot Study

It was immediately apparent in the pilot study that the maternal heart tone signal and maternal aortic pulse wave were saturating the front end amplifier at the gain levels required to detect the fetal tones (gain of 1000 to 10,000). Digital filtering was not feasible because the fetal signal that would be left would be so small that it would disappear in the quantization noise. Two four pole analog high pass filters in series were required in order to eliminate the maternal signals and thereby prevent amplifier saturation. After analyzing spectral data on the fetal and the maternal signal the high pass cut off frequency required was determined to be 20Hz.

The morphology of the recorded fetal heart tone signatures was quite distinct. Figure 2.6 shows a typical recorded fetal heart tone from a patient in the pilot study and clearly depicts the components described in the literature. The mitral, tricuspid, aortic, and pulmonary sub-components are separately apparent, more so than in the examples from the literature. The pilot study revealed that fairly constant heart tone morphology exists across an individual patient's recording, but there is large variability among different patients. This variability is primarily in terms of the relative amplitudes of the four sub-components. There was variation in the number of cycles of a specific component among different patients, but the gross periods of the components were similar. Figure 2.7 compares FPCG recordings from three patients.

The amplitude spectrum of a composite of 20 averaged heart tones from a representative patient is shown in figure 2.8. Note that essentially all of the energy is in the frequency range 5 - 80 Hz with most 15-30 Hz. Composites from other patients in the data collection series show similar spectral plots. These findings are in disagreement

Figure 2.6. FHT Components (from data collection pilot study).

PATIENT 10A

PATIENT 7F

PATIENT 9B

Figure 2.7. Comparison of FPCG Signals from Three Patients.

Figure 2.8. Fetal Heart Tone Spectrum (from data collection pilot study)

with all of the previous literature findings noted above with the exception of Nagel [11]. The data collection pilot study findings show a much lower spectral range for the fetal heart tones than that reported by the investigators other than Nagel, and almost a factor of two lower spectral range than he reported. The marked discrepancies among investigator findings may result from differing responses of the transducer system as well as that of the recording system. The frequency response of the PDVF transducer of this research is essentially flat from 0 - 150Hz (NASA data). The PC data acquisition system also has a flat response in that range. The findings of the data collection pilot study of this research, therefore, would seem to represent the true signature of the fetal heart tone.

The contaminating noise in the raw sensor signal output was found to consist mainly of a large amplitude maternal heart tones and aortic pulse waves (as predicted from the literature), marked maternal body motion artifact, and lesser amplitude environmental sounds. Sixty Hz EMI was also a problem initially due to the high impedance of the sensors. Shielding of the sensors with a grounded copper sheathing (externally insulated to maintain patient electrical isolation) elimnated most of the 60 Hz. To partially solve the anticipated motion artifact problem, the sensors were configured in back-to-back pairs as shown in figure 2.9, which illustrates the seven sensor array front-end electronics. The acoustic wave propagates through the active sensor pair producing signal outputs from each sensor that are 180 degrees out of phase, whereas the motion noise and 60 Hz are in phase. Differentially adding the two outputs cancels a significant portion of the motion artifact and 60 Hz interference while approximately doubling the desired signal output. Unfortunately, the amounts of motion artifact signal on the back-to-back differentially added sensors is not equal because one is in contact with a surface whereas the other is "free." It is unknown whether the relative proportions of motion artifact on the members of a sensor pair is a constant or a function of such factors as belt tightness. Because scaled differential addition is therefore not possible, considerable motion artifact remains to be dealt with by other signal processing methods.
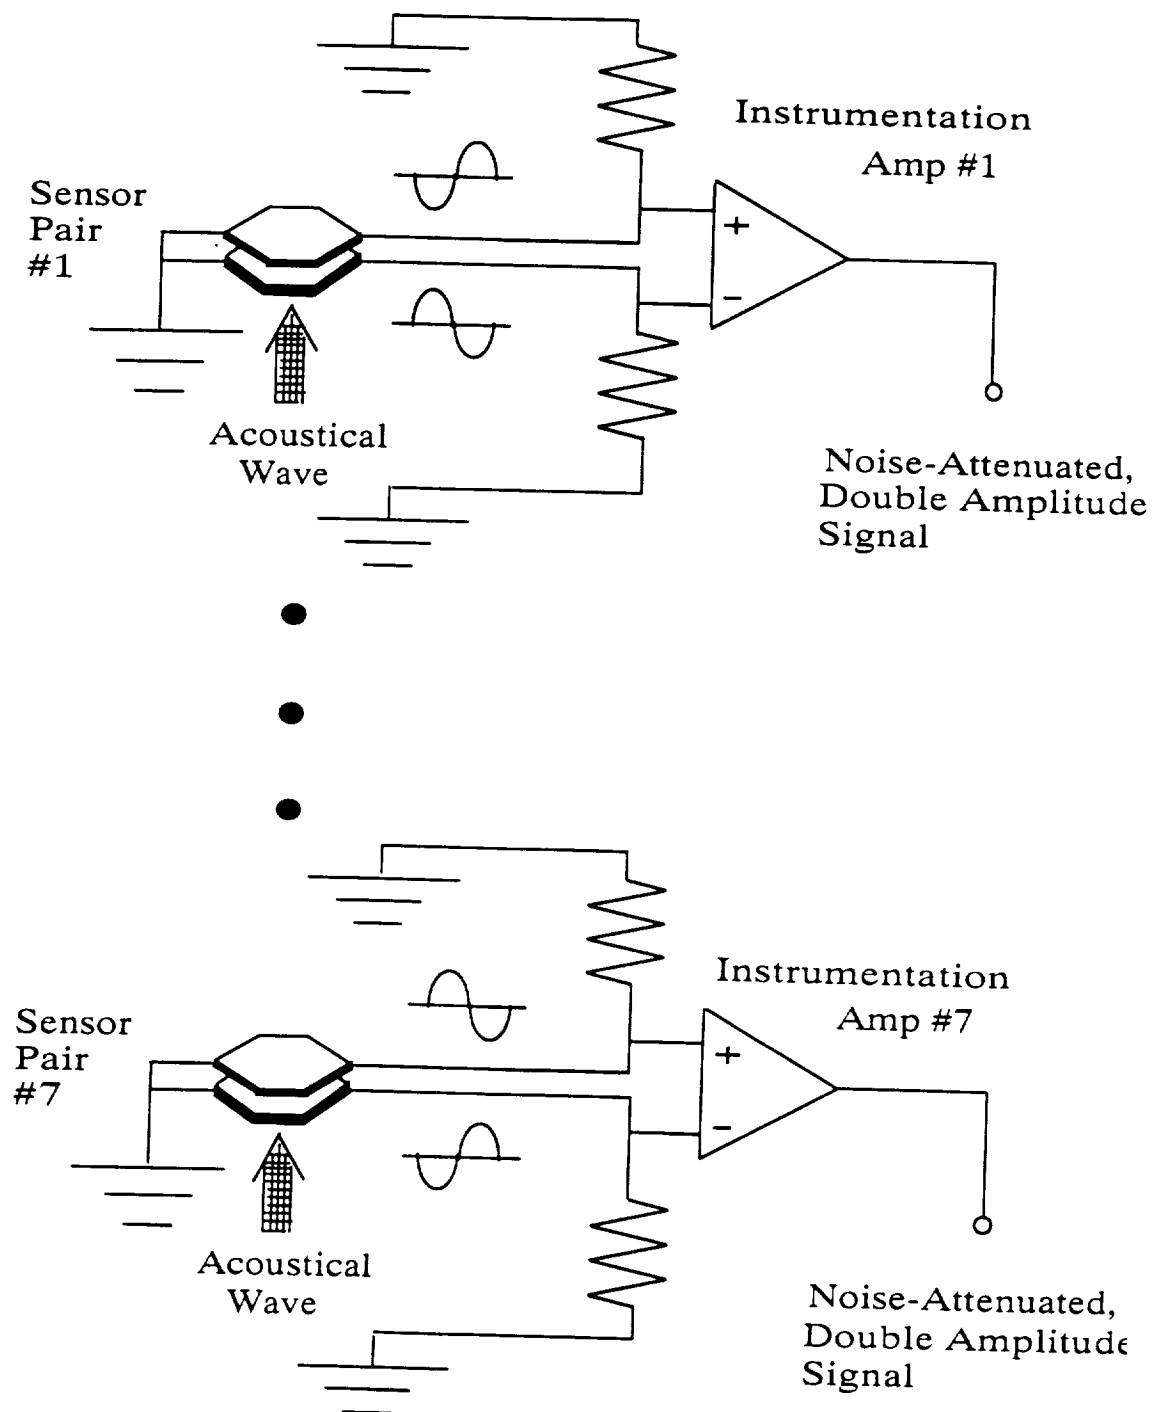
Figure 2.9. Back-to-Back Sensor Design for Differential Noise Cancellation.

In several patients maternal heart tones were also a significant problem in spite of the front-end high pass analog filters. In those patients the maternal and fetal spectrums had large overlap.

2.3 Previous Heart Beat Event Detection Methodologies

Heart beat event (wavelet) detection is necessary in order to determine heart rate. Heart rate is the reciprocal of the time interval between wavelets of the same type (ex. first heart sound of two sequential heart tones). Heart beat wavelet detection methodologies reported in the available literature consist of three approaches: threshold detection, autocorrelation techniques, and signature matching.

2.3.1 Threshold Detection

Threshold detection is the simplest method for wavelet detection and heart rate determination [16]. The point where the wavelet first exceeds a threshold value constitutes the fiducial point for timing. The threshold point is chosen by taking a value for which the probability of exceeding is high for the desired wavelet but which is low for the interwavelet waveforms and/or noise. A minimum signal to noise ratio (SNR) is therefore required for threshold detection, and is

$$SNR > 1.$$

Threshold point values are generally specified relative to the minimum absolute value of the signal. For example,

$$THR = |x_{min}| + 0.2|x_{max} - x_{min}|.$$

There are three disadvantages of threshold detection. First, the overall signal to noise ratio must be significiantly greater than one for the desired wavelets to appear as maxima above the baseline noise level minima. Second, transient noise spikes will be detected as events. Third, baseline drift can cause the detector to fail.

2.3.2 Autocorrelation

Autocorrelation detection is the method of choice for Doppler-shift ultrasonic determination of fetal heart rate [17]. The fetal heart beat is manifested as reflected tone burst wavelets analogous to the fetal phonocardiogram. The tone bursts correspond to Doppler frequency shifts of the reflected continuous outgoing ultrasound signal produced by the moving heart valves. Autocorrelation processing detects the tone bursts among the noise:

$$R_S(k) = (1/N) \sum_{n=0}^{N-1} s(n)s(n-k)$$

where $-N < k < N$. The Doppler shift signal vector block is stored and shifted by k points and the inner product of the stored block with the incoming block produces the autocorrelation function output. Generally, one to three second blocks are autocorrelated in real time to reveal the fetal heart period [17].

Autocorrelation emphasizes periodically occurring correlated wavelets by eliminating non-periodic uncorrelated noise such as movement artifact, baseline drift, and random environmental noise. Autocorrelation processing requires subsequent threshold detection for heart beat event identification.

The disadvantages of autocorrelation processing include its requirement for a periodically recurring wavelet in order for wavelet detection to be accomplished. Variation of the inter-wavelet interval causes a phenomenon called "jitter" when

processing with autocorrelation [13]. Secondly, because the autocorrelation function will emphasize any periodic signal, it will also detect periodic noise[18]. Thirdly, autocorrelation has a high processing time requirement because of the number of multiplications and buffer shifting operations required.

2.3.3 Signature Matching

Detection by signature matching involves comparing the input signal vector to a known desired template vector point by point. This comparison generally is done by either cross-correlation, linear prediction, block transforms, or neural networks.

Cross-correlation:

Cross-correlation of the signal vector with a desired template vector is expressed as follows:

$$R_{xy}(k) = (1/N) \sum_{n=0}^{N-1} x(n)y(n-k).$$

The output of the cross-correlator consists of peaks corresponding to the occurrence of the heart beat wavelets. When the signal block being processed matches with the template, the inner product generates a maximum. In the areas where the signal block does not match with the template, the positive-negative products of the two vectors will tend to cancel out giving a minimum. A threshold detector is subsequently used to detect the peaks [14].

Cross-correlation of the signal with a template has the advantage of detection of frequency sequential relationships, i.e. the heart wavelet signature. Disadvantages are its sensitivity to phase shifts and the high number of mutiplications and buffer shifts required.

Linear Prediction:

Linear prediction is mathematically similar to cross-correlation. Linear prediction theory states that if a signal signature represents the response of a linear system, then each point in the signal signature vector is a linear combination of all the preceding points, or:

$$\hat{S}(n) = \sum_{k=1}^{N} S(n-k)P(k)$$

where $S(n)$ are the signal vector samples and $P(k)$ are the predictor coefficients. ECG signals, for example, are considered predictible through second order auto-regressive modeling [19]. Lin and Chang [20] used a linear predictor, also of order two, to detect ECG wavelets in adults. The residual prediction error,

$$E(i) = \hat{S}(i) - S(i)$$

reaches local minimums at the points of the occurrence of the ECG wavelets in the raw signal. A simple error threshold is then used as the basis of a wavelet detector. They furthermore found that increasing the prediction order above two does not affect the error between the predicted and actual signal.

Lin and Chang computed the predictor weights by means of a recursive procedure carried out on a "short-time" autocorrelation coefficient matrix derived from heart signal templates. A mean square minimization technique was used for optimization of the weights. The authors also point out that, although in speech analysis the higher the order the better the predictor performance, such does not appear to be true for ECG wavelet

detection. The advantages of the linear prediction detector as compared to alternative methodologies include more accurate detection of signal features and higher speed.

Block Transform Methods

Taking the block transform (ex. FFT) of the incoming signal vector generates its spectral parameters. Directly comparing these spectral parameters to those of a template heart tone signature constitutes frequency domain detection. Venkat [21] used a method for speech processing of taking a second FFT of the initial FFT spectral parameters to generate a set of four or five identifying "features." These features are then used to map the block transform result to a feature space. There a probability decision is made by means of a Euclidean distance measurement versus a distance threshold as to whether a desired waveform has been detected. Such transform computations require considerable processing time and may not detect an actual signature but rather simply its frequency make up.

Neural Nets

Training a neural network with a known heart signal template is a further method used for heart signal signature detection. Neural net processing, however, tends to work best in offline waveform processing such as for ECG analysis and has not been found practical for real time heart beat event detection.

2.4 Detection Methodology for this Research

The detection of the fetal heart tones is not a simple task. The low baseline SNR due to ambient background noise and the large-amplitude maternal heart tones, aortic pulse waves, placental flow sounds, and bodily movement artifact require a highly specific as well as a highly sensitive detection process.

2.4.1 Summary of Detection Methodologies (Pros and Cons)

The advantages and disadvantages of the above detection methods as applied to this research are as follows:

Simple threshold detection requires a SNR higher than is expected in the fetal signal. Furthermore, threshold detection would interpret noise spikes such as maternal movement and bodily sounds as fetal heart beat events.

Autocorrelation would not be an optimum technique for detection because of its tendency to lock on other periodic noise such as 60 Hz EMI or maternal heart signals. Autocorrelation would produce jitter because the fetal heart beat is not exactly periodic.

Block transforms and neural nets suffer from time constraints imposed by real time applications. Block transform detection may also tend to lock on collections of frequencies similar to those of the fetal heart tones rather than the actual signature of the heart tones.

Template cross-correlation matching would theoretically detect the actual heart tone signature. Although it is mathematically similar to linear prediction, cross-correlation is less precise than linear prediction because cross-correlation does not detect actual inter-relationships between signal points.

Linear predictors are computationally efficient, they work in low SNR's, and they detect an actual signature because they are sensitive to inter-relationships between signal points. The linear predictor was therefore chosen as the fetal heart tone detector for this research.

2.4.2 LMS Linear Prediction Algorithm

The expression for the predictor operation again is:

$$\hat{X}(i) = \sum_{k=1}^{L} X(i-k)W(k)$$

where X(i) is the actual current sample value and $\hat{X}(i)$ the value predicted from past sample values, W(k) is the weight vector, and L is the length of the predictor. The predictor error is

$$E(i) = \hat{X}(i) - X(i).$$

The error value is used in two ways. In the "training" mode the error is used in a feedback loop to adjust the values of the weights to more accurately predict the desired signal vector values. The weights are iteratively adjusted until the mean square error is minimized. At that point the weights optimumly describe the system generating the signal or have "modeled" the system [14]. In the processing mode the weights are constants, having been derived in the training mode. When the mean square error in the processing mode reaches a local minimum, the predictor has "locked onto" the signal from which it's weights were derived and detection is accomplished.

The number of weights required to accurately model a particular system response (the "order" of the predictor) is system dependent. For example, between five and 22 weights are required to describe the adult second heart sound system response [22] whereas only two weights are required to predict ECG [19].

The mean of the squared predictor error (MSE) can be shown to be a quadratic function of the weights [23]. If this quadratic function is plotted versus MSE with N=2 weights, for example, the plot will form a three dimensional bowl called a "performance surface." The minimum point on this performance surface is the minimum MSE and is found by taking the gradient of the equation of the surface. With multiple weights a multidimensional performance surface is formed, but the minimum MSE is determined by the same gradient method. This gradient approach produces an expression for updating the weights in predictor training by minimizing the MSE called the steepest

descent method. The expression is, however, untenable for use on a computer. A compromise approach to the steepest descent method is called the "noisy" approximation to gradient estimation. The noisy gradient weight update expression is

$$W(k+1) = W(k) + 2\sigma E(k)X(k)$$

where σ is a constant regulating the stability and speed of convergence of the adaptation. When the weights converge at the minimum MSE value, the predictor has modeled the desired signal.

The above weight adjustment operation constitutes the Widrow-Hopf Least Mean Square (LMS) algorithm [23][14]. The LMS algorithm is the most commmonly used algorithm for adaptive filters. A diagram of the LMS algorithm is shown in figure 2.10 which illustrates the predictor training and processing modes.

2.4.3 Training and Generation of Predictor Weights

The predictor weights are generated via the Widrow-Hopf LMS Algorithm by training on the desired signal signature, i. e. a template. A template is an average or generic wavelet. The training process involves multiple passes over the template until convergence of the mean square predictor error occurs.

A template can be artificially synthesized or generated from real data. There are two methods for real data template generation: time domain wavelet ensemble averaging and frequency domain spectral averaging with time domain reconstruction.

Time Domain Template Generation:

This method assumes that the heart tone signal is periodic and ergotic. The noise is assumed to be ergotic and wide-sense stationary. One can therefore obtain an ensemble average by averaging sequential time wavelets. If the noise is not periodic and
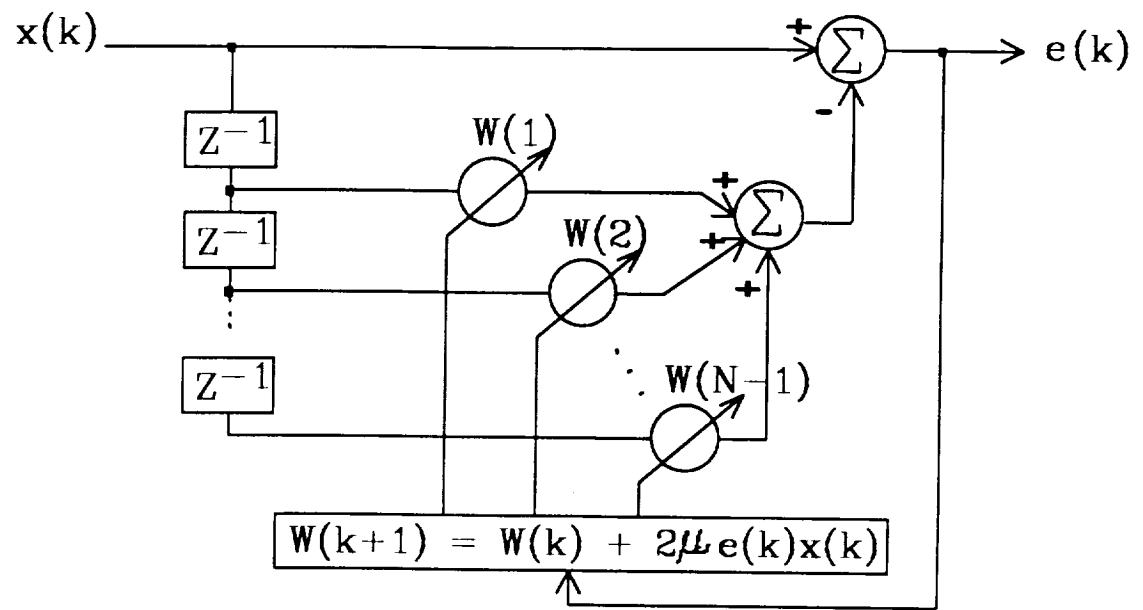
(a)

(b)

Figure 2.10. Linear Predictor Block Diagram:
a) Training mode with LMS Algorithm. b) Processing mode

has zero mean, then averaging of multiple wavelets in the time domain will tend to cancel the noise. As the number of wavelets averaged goes to infinity the noise term tends to go to zero by cancellation of the random positive and negative deflections .

Aligning wavelet blocks for ensemble averaging is a problem. If the wavelet samples to be averaged are not in the same relative time positions from one wavelet to the next, then an inaccurate template will result. Defining boundaries by uniform time blocking starting at some arbitrary fiducial point results in jitter because of the non-exact periodicity of the heart wavelets [13].

Frequency Domain Template Generation:

Template generation by frequency domain averaging is less complicated than time domain averaging. Synchronization of wavelet blocks is not necessary. Fast Fourier Transforms (FFT) are taken of wavelet blocks of uniform length but with the wavelet having an arbitrary position within the block. The FFT's of multiple wavelet blocks are then averaged in terms of the individual frequency components. The inverse is then taken of the average FFT to give the wavelet template.

As with time domain averaging, there are problems with template generation by frequency domain averaging. Time domain averaging cancels random background noise and uncorrelated heart sounds, but alignment of the wavelet points is a problem. Frequency domain averaging involves minimal alignment problems but background noise is not automatically elimated. Even if background noise is random, if the frequencies are uniform throughout the signal, the noise will be represented in the template [13].

Initially in this research the template for predictor training was generated in the frequency domain by averaging FFT's of multiple representative heart tone blocks (usually twenty to fifty) and then taking an inverse FFT of the average. Because of

inadequate detector performance with the frequency domain method, a time domain training method was implemented. But unlike the above described time domain method, which requires alignment and averaging of time wavelets, the method used by this research involves training on whole patient files. Whole file training consists of passing the training predictor over all the heart tones in a patient file, consisting of up to 2000 heart tones. As the file is processed by the training algorithm the weight update routine is turned on only when a heart tone is encountered and turned off between tones. The turn on point, T_E, is governed by a rise in the local signal power above an arbitrary threshold, here set to be:

$$T_E = E_{min} + Q [E_{max} - E_{min}].$$

Where Q was empirically set at about 0.2. One problem with this energy threshold method is that noise spikes can artificially boost E_{max}. Therefore, E_{max} was set at the nominal maximum observed heart tone energy level. Also, the weight updating was turned off when the local energy level exceeded E_{max}. These adjustments produced better training results for the whole file method.

A comparison of frequency domain template versus time domain whole file training was carried out on several patient files. The whole file training method gives superior results. Accordingly, the whole file method was used to generate weights for real time implementation of the LMS linear predictor algorithm.

Weights generated from one patient file were found to be suitable for almost all other patients, indicating global signature characteristics. An optimum set of weights derived from a single patient in the data collection phase was used for all patients in the real time NST comparison trials described in Chapter 4.

2.4.4 Linear Predictor DC Offset Sensitivity Correction

It was observed in the initial real time implementation of the LMS linear predictor algorithm that small changes in the DC bias of the signal produced disproportionally large changes in the predictor error. The predictor error shift is constant if the DC level is constant, which is not significant if the DC level is low relative to the signal amplitude. But there is also the local DC level generated as a result of processing finite length signal blocks . The predictor error shift from this local DC is continually changing and therefore results in a noisy predictor error. The reason for the DC effect and a way to correct it can be understood from the following analysis. The predictor equation with the DC term included is

$$\hat{X}(n) = \sum_{k=1}^{N} W(k)[X(n-k) + C]$$

where C is the mean value of the signal block. Expanding this expression:

$$\hat{X}(n) = \sum_{k=1}^{N} W(k)X(n-k) + C \sum_{k=1}^{N} W(k).$$

Thus, the amount of the deviation of the predictor result is the inner product of the DC value with the predictor weights, which explains the disproportionate effect of DC bias on the predictor result. One does not want to totally eliminate local signal DC bias from the predictor result since this is part of the signature one is trying to detect. The predictor result that is desired is

$$\hat{X}(n) = \sum_{k=1}^{N} W(k)X(n-k) + C$$

Therefore, subtract $C \sum_{k=1}^{N} W(k)$ and add C:

$$\hat{X}(n) = \sum_{k=1}^{N} W(k)X(n-k) + C \sum_{k=1}^{N} W(k) - C \sum_{k=1}^{N} W(k) + C$$

Now reorganizing to obtain the desired correction term,

$$\hat{X}(n) = \sum_{k=1}^{N} W(k)X(n-k) + C \sum_{k=1}^{N} W(k) - C [1 - \sum_{k=1}^{N} W(k)]$$

Remembering that C is the signal vector mean,

$$\hat{X}(n) = \sum_{k=1}^{N} W(k)X(n-k) + C \sum_{k=1}^{N} W(k) + [1/N \sum_{k=1}^{N} X(n-k)] [1 - \sum_{k=1}^{N} W(k)]$$

$$\hat{X}(n) = C \sum_{k=1}^{N} W(k) + \sum_{k=1}^{N} X(n-k)[W(k) + [1 - \sum_{k=1}^{N} W(k)] / N]$$

Thus, to eliminate perturbations due to DC bias and the local signal block mean, the factor $[1 - \sum_{k=1}^{N} W(k)] / N]$ is added to each of the weights ahead of time.

The above correction factor was implemented in the predictor training procedure and results in a significant improvement in the predictor performance as evaluated in a pilot study.

CHAPTER THREE

REAL TIME IMPLEMENTATION

3.1 Introduction

The methodologies are presented in this chapter for the real time implementation of the LMS linear prediction algorithm for the detection of fetal heart tones. The technique used for heart rate derivation and correction for spuriously detected heart tones is explained. In section 3.2 the hardware system is described. An overview of the software is given in Section 3.3. The TMS320C25 real time algorithm is presented in Section 3.4. The PC/TMS handshaking control/display program is delineated in section 3.5. The algorithm variables are listed in section 3.6. Processor utilization is discussed in section 3.7.

3.2 Hardware

A block diagram of the monitor system hardware is shown in figure 3.1. The hardware consists of the seven-element sensor belt, an impedance matching instrumentation amplifier for each sensor with a fixed gain of 100, an analog multiplexer to allow sensor selection, two four pole Butterworth 20 Hz high pass filters in series, a six pole Bessel 55 Hz anti-aliasing low pass filter, a second amplifier with gain adjustable from one to 500, a 40 MHz TMS320C25 (TMS) digital signal processing board (Atlanta Signal Processing, Inc.), a PC, and a single channel analog strip chart recorder (AB Goerz Co., Model SE110). Additionally, a handheld button allows patients to flag the occurrence of fetal movement on the strip chart tracing. The hardware is

Figure 3.1. Hardware Block Diagram.

housed in a portable cart so that it can be easily transported to the EVMS clinic and fit into the tight confines of the patient evaluation rooms. Figures 3.2 ,and 3.3 depict the hardware system.

The hardware parameters were determined by analysis of the information gathered from the data collection phase. The real time sampling rate was selected to be 158 Hz because greater than 90% of the fetal heart tone spectral energy was noted to be below 55 Hz (figure 2.8) and because 158 Hz was the lowest sampling rate available for the TMS320C25 board. The anti-aliasing cutoff frequency of 55 Hz was selected to be as low as possible without significantly attenuating the fetal signal. A sharp 55 Hz cutoff also would attenuate some of the 60 Hz EMI.

3.3 Software Overview

The system software consists of the TMS software and the PC software which run concurrently in real time.

The TMS software processes the filtered acoustic sensor signal for the detection of fetal heart tones. The TMS software also accomplishes heart rate derivation, sensor selection, and calibration and control of the strip chart recorder.

The PC software is a PC/TMS control/display routine. Its function via handshaking with the TMS is to allow PC keyboard control of the TMS operation and to display in real time the incoming acoustic waveforms along with the parameters generated by the TMS algorithm and the output heart rate value. Also displayed are an accumulator overload flag and a poor signal quality indicator.

3.4 TMS320C25 Algorithm

An overview flow diagram of the TMS Fetal Heart Rate Monitor Algorithm is shown in figure 3.4. The TMS software includes the Initialization Routine, the Linear

Figure 5. Hardware setup for Portable Case
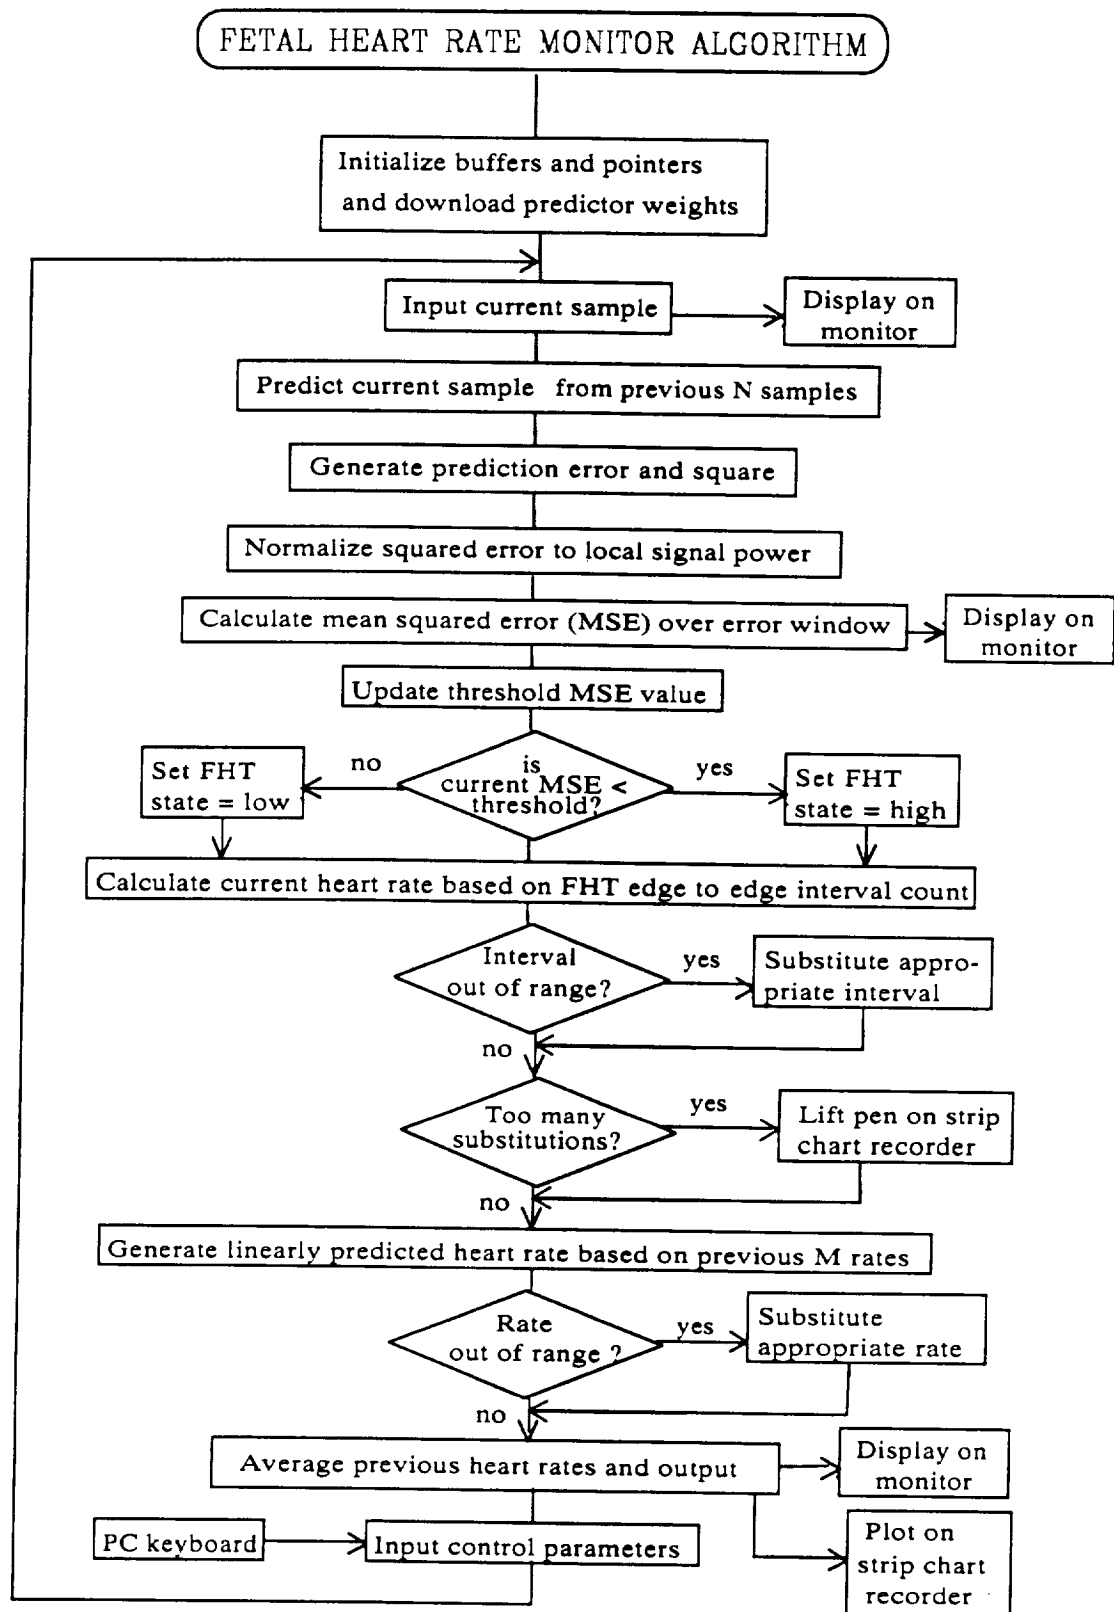
Figure 3.4. Real Time Fetal Heart Rate Agorithm Overview

Predictor Routine, the Normalization Routine, the Mean Square Error (MSE) Routine, the Threshold Routine, the Heart Rate Routine, and the Signal Quality Routine. After the Initialization Routine a main loop routine calls each of the other routines as discussed below in section 3.4.2.

3.4.1 Initialization Routine

On start up the Initialization Routine, shown in figure 3.5, initializes several buffers and pointers. The weights for the heart tone linear predictor and the heart rate linear predictor are down-loaded as well as several constants.

3.4.2 Main Loop Routine

The Main Loop Routine shown in figure 3.6 constitutes the core of the TMS software. It is entered after initialization and calls all of the other routines in sequence. Control parameters are lastly input and the loop repeats.

The Main Loop Routine begins with the input of a sample via the TMS board A/D converter. Start of conversion is initiated by the on-board sampling rate clock. The end of conversion flag is polled by the routine until activated at which point the current sample is input and stored. The Linear Predictor Routine is then called.

3.4.3 Linear Predictor Routine

The Linear Predictor Routine is shown in figure 3.7. The predictor estimates the current sample point by taking an inner product of the immediately preceding N samples with the predictor weights. The predicted value is defined by

$$\hat{X}(n) = \sum_{k=1}^{N} X(n-k)W(k)$$
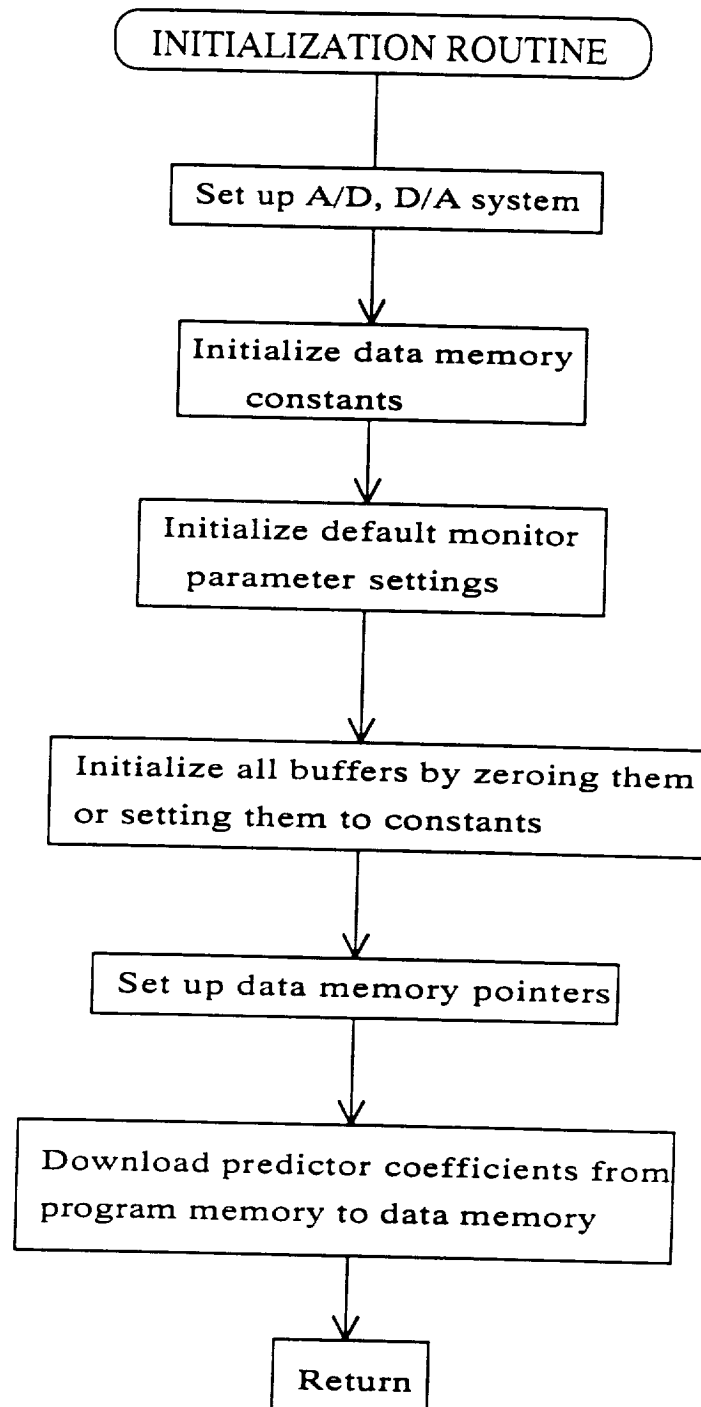
Figure 3.5. Initialization Routine.
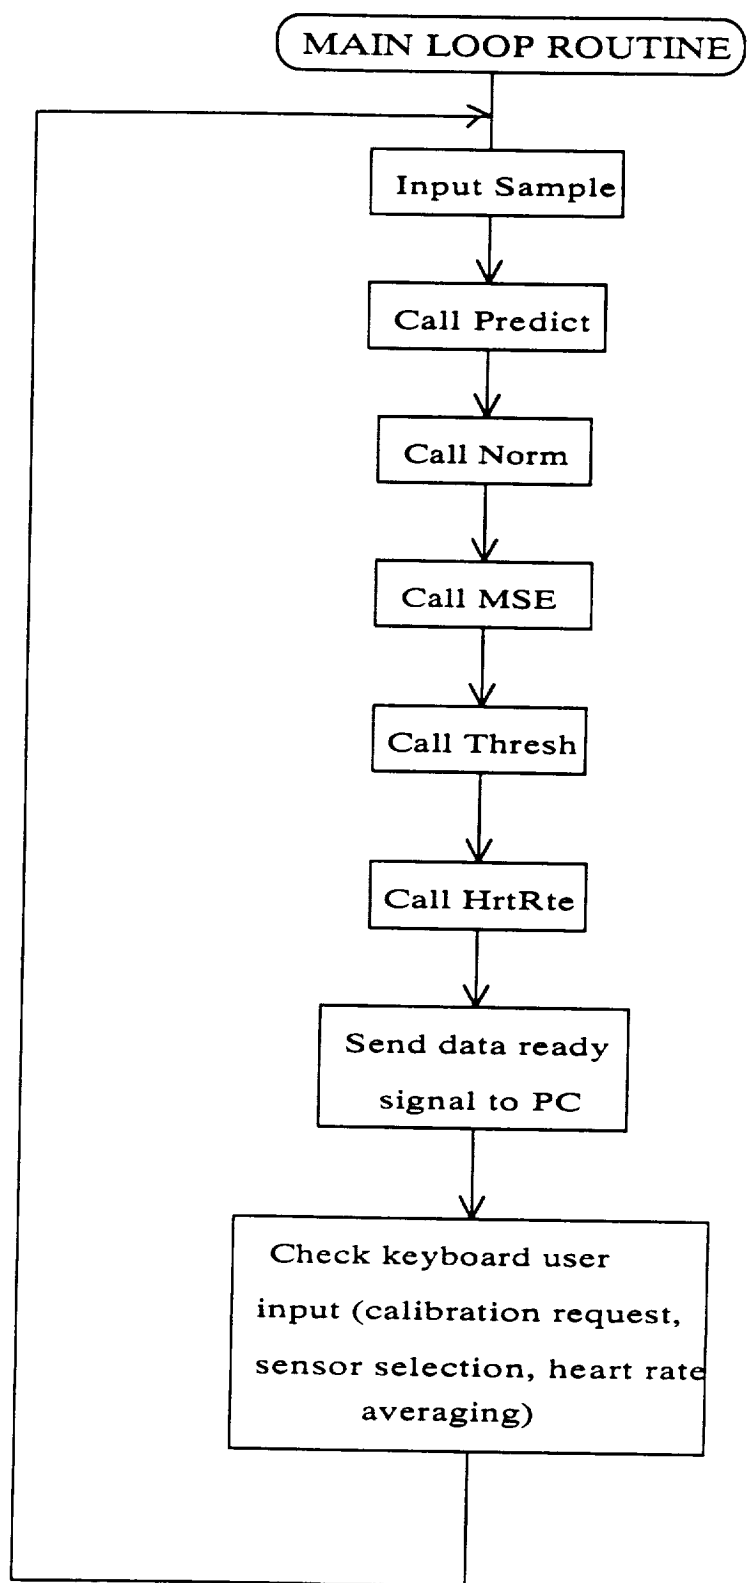
Figure 3.6. Main Loop Routine.
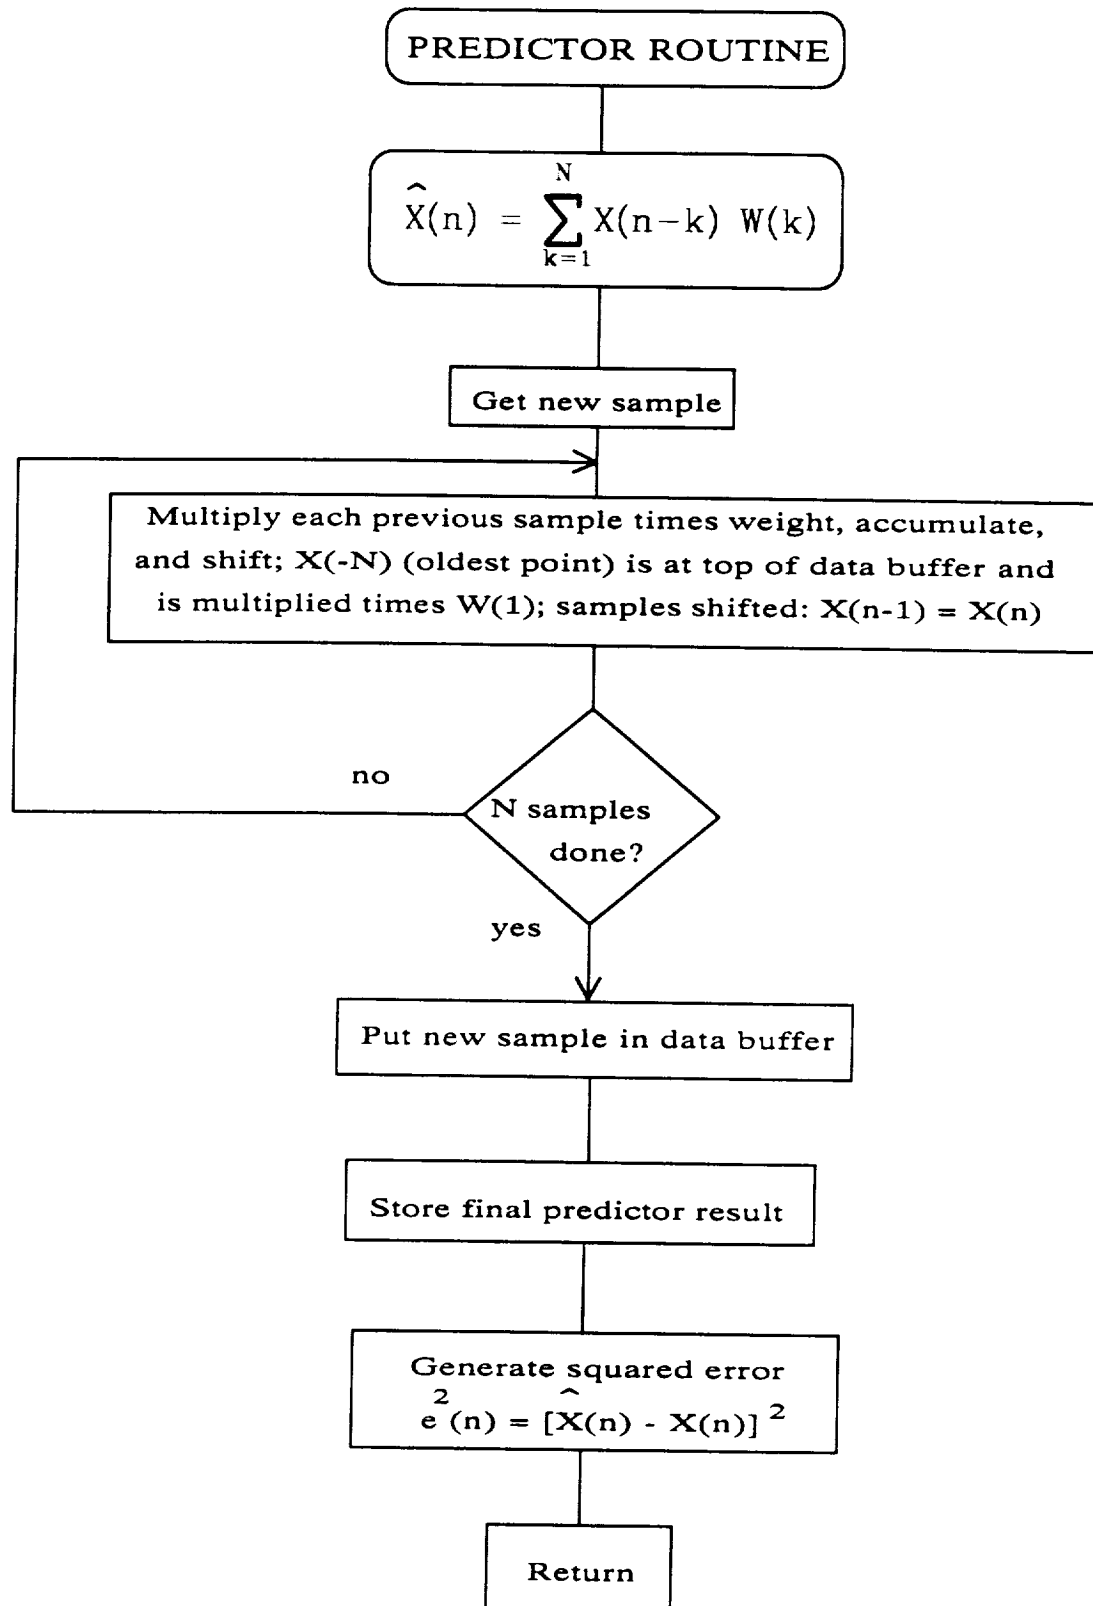
Figure 3.7. **Linear Predictor Routine.**

At the same time the samples in the data buffer are rotated (shifted upward by one location, the oldest sample being lost) and the current sample is placed at the bottom of the buffer. Fixed point arithmetic on the TMS requires that the values in the inner product be integers and that the inner product result not overflow the 32 bit accumulator, including the sign bit. The sample values, which come from the A/D converter, are already 16 bit integers including the sign bit. The predictor weights, which are generated off-line, are decimal values of slightly less than 2 and so are scaled for TMS arithmetic as 16 bit integers (including sign bit) by multiplying each by 2 14. The product of one weight and one sample could be maximally 32 bits. But N 32 bit numbers summed for the inner product could overflow the accumulator when the input signal is of large amplitude. The accumulator overflow flag is therefore output to the PC monitor display as an indicator of the need to decrease front end amplifier gain. As well as observing the overflow flag, the user can also adjust the amplifier gain so that the waveform on the display is within a specified range of amplitude. Automatic gain control was not used due to complex scaling problems encountered with implementation of this process in fixed point arithmetic.

The final 32 bit inner product predictor result is converted to a 16 bit signed result by shifting the accumulator two places to the left and storing the high 16 bits. This divides the result by 2^{14} to compensate for the initial scaling up of the weights by 2^{14}.

The predicted sample value is subtracted from the actual sample value encountered to generate the predictor error which is squared in preparation for subsequent generation of the mean square error expressed by

$$e^2(n) = [\ X(n) - \hat{X}(n)\ ]^2$$

where $e^2(n)$ is the squared predictor error. The squared error is then normalized to the local signal power in the Normalization Routine.

3.4.4 Normalization Routine

The Normalization Routine is shown in figure 3.8. Normalization of the squared error to the local mean signal power is necessary because the difference between the predicted signal point and the actual signal point is a not only a function of how good a match the predictor has made but also is a function of the amplitude of the signal itself. The latter is due to the fact that the predictor inner product contains the signal vector, so that the above difference value is a scalar function of the signal amplitude. The error would therefore tend to increase when a heart tone was encountered rather than decrease if normalization to local mean signal power were not performed. Mean signal power is used for normalization rather than mean amplitude because it is the square of the error that is being normalized.

The local mean signal power estimation is accomplished by summing the squares of the signal sample amplitudes over the predictor block length. DC bias should not be included in the local power estimation so that the variance of the signal vector is used. The expression for the normalized squared predictor error is then

$$\tilde{e}^2(n) = \frac{e^2(n)}{\sigma_x^2}$$

where $\tilde{e}^2(n)$ is the normalized squared error and σ_x^2 is the variance. An unbiased estimator for the discrete point variance is

$$\text{NORMALIZATION ROUTINE}$$

$$\tilde{e}^2(n) = \frac{e^2(n)}{\sigma_x^2} = \frac{e^2(n)}{\sum_{i=0}^{N-1} X^2(i) - \frac{1}{N}\left[\sum_{i=0}^{N-1} X(i)\right]^2}$$

To obtain numerator sum signal points coming in; multiply times scaled decimal 1/N and rescale; subtract result from sum of squared signal points

To accomplish division, determine number of leading zeros of quantities in numerator and denominator (32 bit quantities)

Call Leading Zero Routine

Subract numerator leading zeros from denominator leading zeros ==> power of 2 normalized sq. error

Return

Figure 3.8. Normalization Routine.

$$\sigma_X^2 = [1/(N-1)] \sum_{n=0}^{N-1} [X(n) - \bar{X}]^2 .$$

Implementation of the above expression on the TMS requires breaking it apart and then recombining to produce the desired form. Thus,

$$\sigma_X^2 = [1/(N-1)] \sum_{n=0}^{N-1} [X^2(n) - 2\bar{X}X(n) + \bar{X}^2]$$

$$= [1/(N-1)] [\sum_{n=0}^{N-1} X^2(n) - 2\bar{X} \sum_{n=0}^{N-1} X(n) + \sum_{n=0}^{N-1} \bar{X}^2]$$

$$= [1/(N-1)] [\sum_{n=0}^{N-1} X^2(n) - 2N\bar{X}^2 + N\bar{X}^2]$$

$$= [1/(N-1)] [\sum_{n=0}^{N-1} X^2(n) - N\bar{X}^2]$$

$$= [1/(N-1)] [\sum_{n=0}^{N-1} X^2(n) - N[(1/N) \sum_{n=0}^{N-1} X(n)]^2]$$

$$\sigma_X^2 = [1/(N-1)] [\sum_{n=0}^{N-1} X^2(n) - (1/N)[\sum_{n=0}^{N-1} X(n)]^2] .$$

The expression for the normalized squared error therefore becomes:

$$\tilde{e}^2(n) = \frac{[e^2(n)] (N-1)}{\sum_{n=0}^{N-1} X^2(n) - (1/N)[\sum_{n=0}^{N-1} X(n)]^2}.$$

The factor N - 1 in the numerator can be eliminated since it is simply a constant, so that the final expression for the normalized squared error is:

$$\tilde{e}^2(n) = \frac{e^2(n)}{\sum_{n=0}^{N-1} X^2(n) - (1/N)[\sum_{n=0}^{N-1} X(n)]^2}.$$

The denominator quantities are now easily computed on the TMS. The sum of the squares of the signal samples and the square of the sum of the signal samples are accomplished via separate rotating buffers as the samples come in. To streamline division by N, 1/ N is represented as a scaled integer and the computation is accomplished by multiplication and subsequent rescaling which is more time efficient for the TMS.

Implementation of the final expression on the TMS requires division of the squared error by the expression in the denominator. This creates a significant problem. Division on the TMS, is a relatively complex, time-consuming process, and is especially so for 32 bit division. Thirty-two bit division is necessary in order to avoid untenable scaling problems and sub-optimal truncation of accumulator results, as each of the above denominator separate sums can be much larger than 16 bits. A method that would consistently grab the optimum 16 bits of the 32 bit accumulator sums and thereby allow 16 bit division for error normalization was unable to be formulated. Also, fixed point division requires a separate routine for integer versus fractional division which again

results in considerable scaling problems when manipulating quotients that may come from the two different routines.

In order to circumvent the above difficulties, 32 bit division is implemented by a leading zero comparison routine (figure 3.9). First, the number of leading zeros of the numerator and denominator quantities are determined by a shift into carry, branch on carry methodology. Then the number of leading zeros of the numerator is subtracted from that of the denominator to give an exponential normalized squared error, NSE_{exp} or

$$NSE_{exp} = \log_2 (\tilde{e}^2(n)).$$

3.4.5   MSE Routine

The MSE Routine, as shown in figure 3.10, converts the exponential normalized squared error into a real numerical squared error and then computes the mean squared normalized error (MSE). Mathematically this is

$$MSE(n) = (1/E) \sum_{n=0}^{E-1} (2)^{NSE_{exp}}.$$

where E is the error window. This conversion is accomplished by loading the accumulator with all zeros except for a "1" in a single bit position. The accumulator is rotated the number of times equivalent to the value of NSE_{exp}, left or right depending on the sign of the exponent, or not at all if the exponent is zero. The optimal starting bit position for the shifted "1" was determined experimentally by observation of NSE_{exp} ranges during real time fetal heart tone processing. The resulting 32 bit normalized squared error is equivalent to 32 bit fixed point division normalization and is produced with no scaling problems.
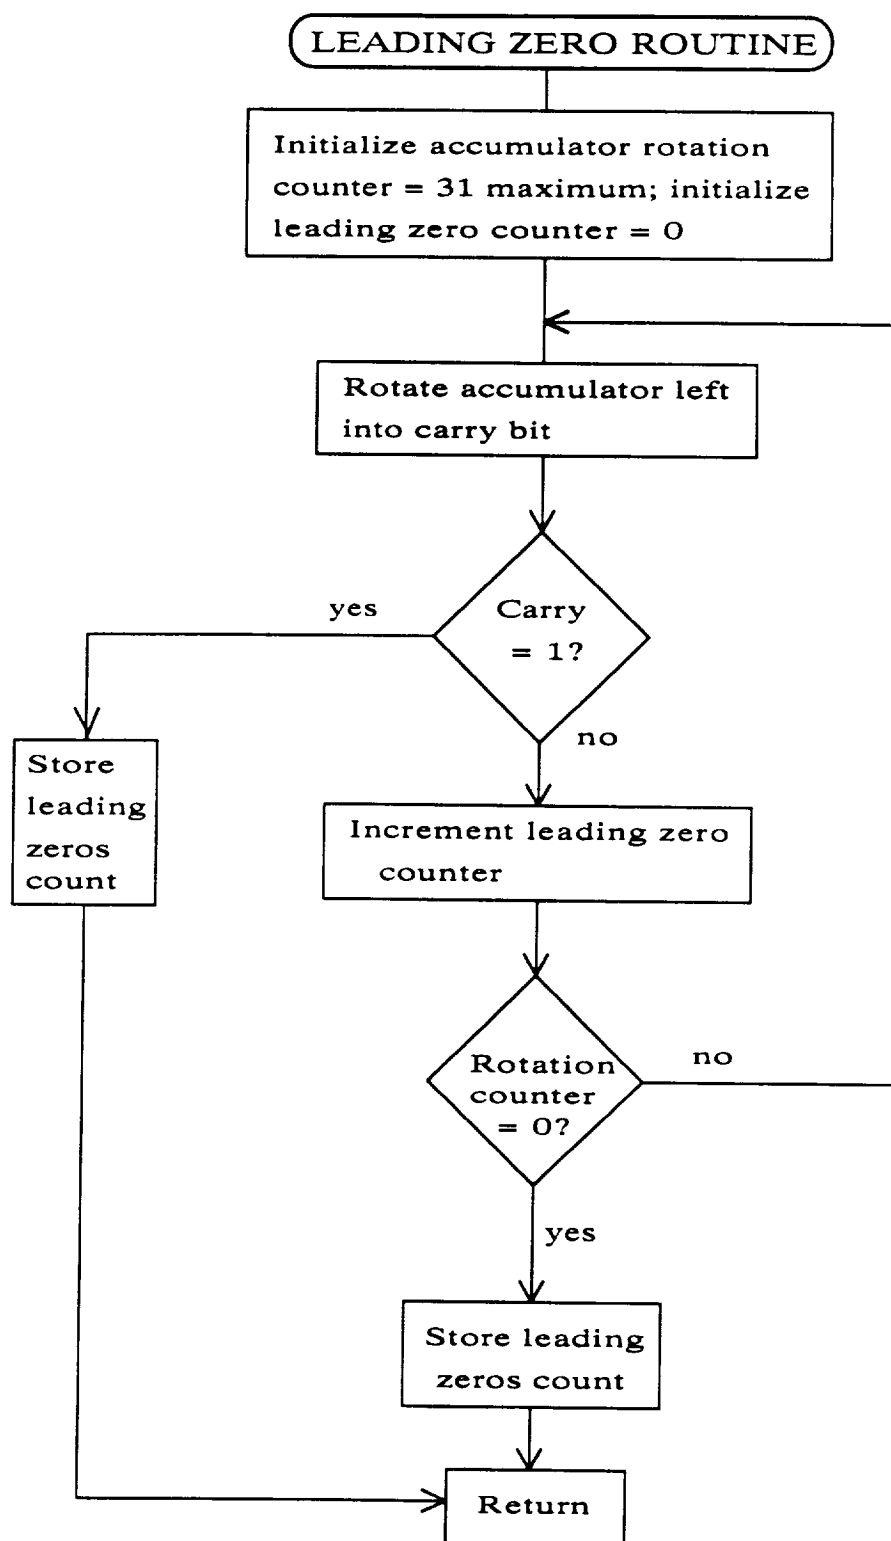
Figure 3.9. Leading Zero Routine.
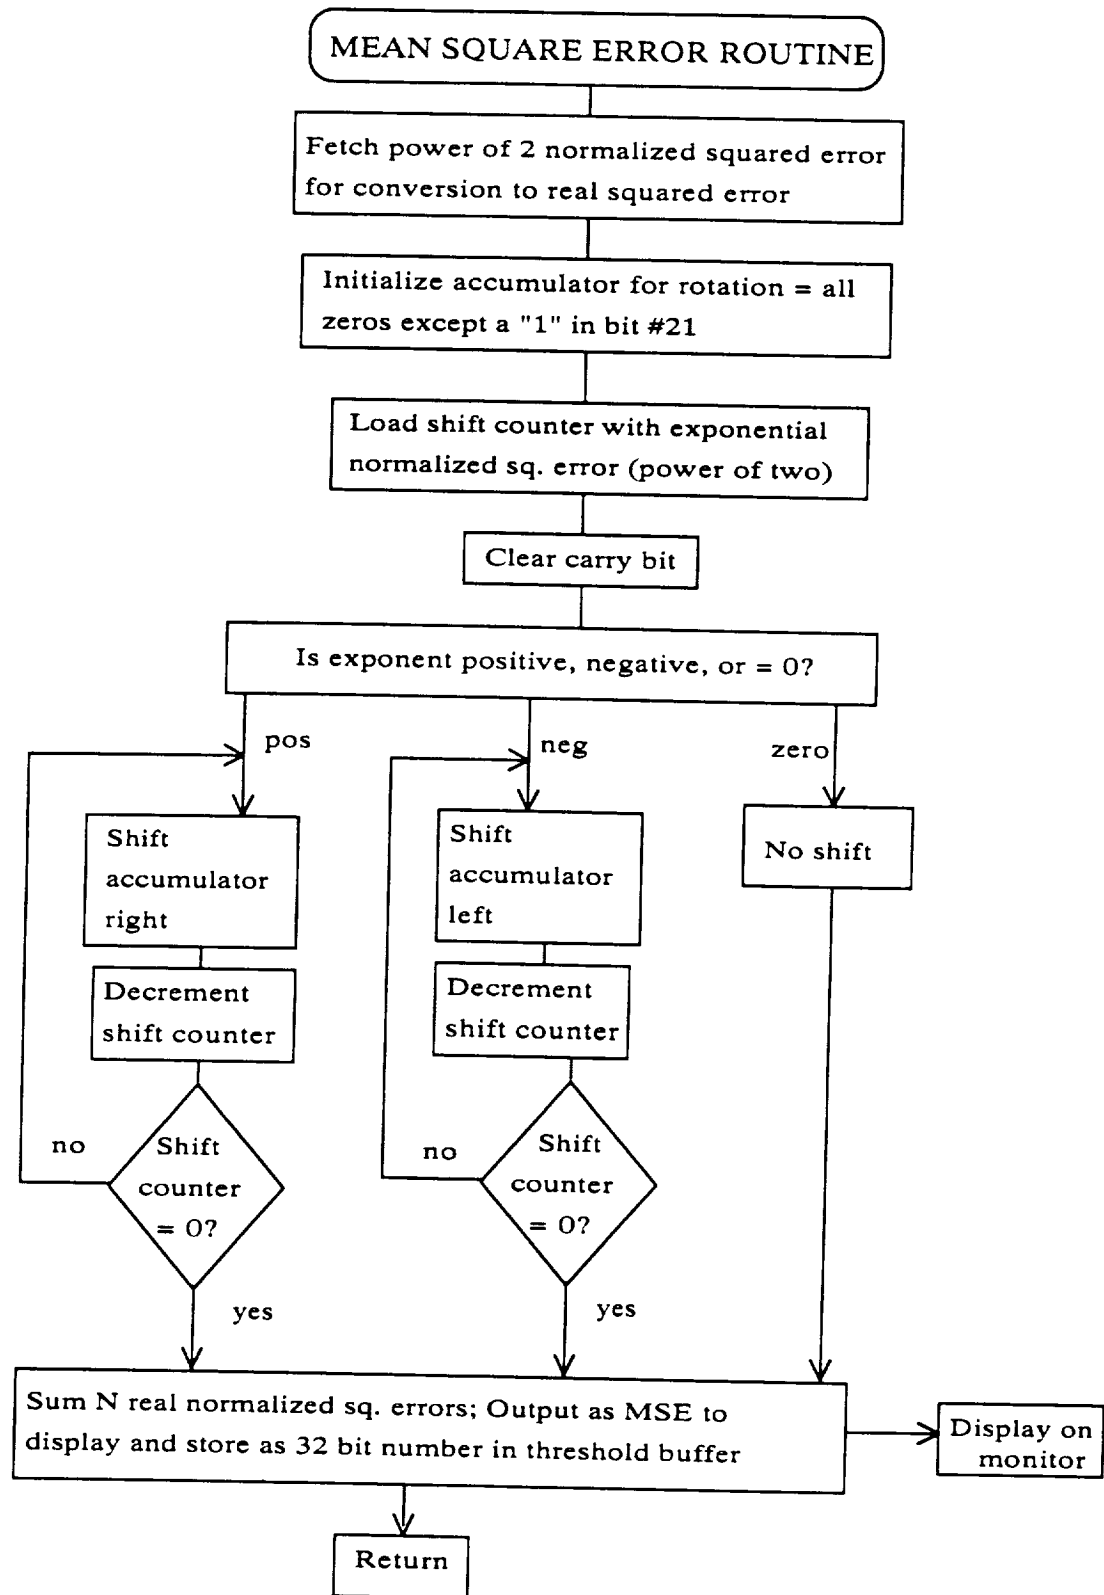
Figure 3.10. Mean Square Error Routine.

This leading zero method of division gives a result that is theoretically accurate to within a factor of two. Confirmation of this theoretical accuracy was accomplished by a pilot study. In this pilot study the fixed point values were calculated by a PC on the above numerator and denominator quantities transferred in real time from the TMS. The corresponding leading zero normalized squared error was also transferred from the TMS for comparison. The maximum difference between the leading zero fixed point values and the floating point values was observed to be a factor of two or less.

The remainder of the MSE routine is devoted to calculating the mean of the current and previous normalized squared errors over an error window. A rotating buffer was used for storage of the 32 bit normalized squared error values and calculation of the MSE. The length of the error window was arbitrarily set to be the length of the predictor.

The 32 bit MSE value is read from the TMS memory by the PC and plotted on the monitor screen in real time underneath the fetal signal as illustrated in figure 3.11 (a synthesized FPCG signal is produced by repeatedly outputting a template set of points). The level of the MSE ostensibly indicates the presence or absence of the fetal heart tone signature. The point of decline of the MSE lags the fetal heart tone plot by 2N points due to the length of the predictor and the length of the error window.

When the MSE value declines to a local minimum, the predictor has locked onto a signature similar to that from which its coefficients were derived. Noise waveforms give local minimums theoretically of less magnitude, so that by setting a minimum MSE threshold, detection of the fetal heart tone signature is accomplished.

Figure 3.11. Real Time PC Monitor Display (synthetically generated heart tones by repeating template)

3.4.6 Threshold Routine

The Threshold Routine is illustrated in figure 3.12. The mean of the most recent block of MSE's is used as the threshold, $H_{MSE}(n)$, for detection of the fetal heart tone signature. This is expressed by

$$H_{MSE}(n) = (1/M) \sum_{k=0}^{M-1} MSE(k)$$

where M is the threshold calculation window. The sample index number, n, is included in the expression because the threshold value is updated with each new sample value. A continually updating threshold level was used instead of a global threshold level in order to accommodate changes in signal quality. Several threshold window lengths were evaluated in a pilot study in order to determine an optimal range, with M = 128 being used in this research (for division convenience M is chosen as a power of two).

The threshold routine uses a pair of rotating 16 bit buffers of length M to accomplish 32 bit arithmetic. Each new 32 bit MSE value is stored as two 16 bit numbers placed at the bottom of the buffers. The values in the buffer are then summed and rotated (shifted to the next higher locations with the oldest value being lost). The mean of the M summed MSE's is taken by shifting the accumulator right by $log_2 M$ bits and storing the low 16 bits.

The MSE threshold level is next combined with positive and negative hysteresis values. This hysteresis-modified result constitutes the final threshold (range) for comparison with the current MSE for fetal heart tone (FHT) detection. The amount of hysteresis ideally should be the standard deviation of the background noise, but as this
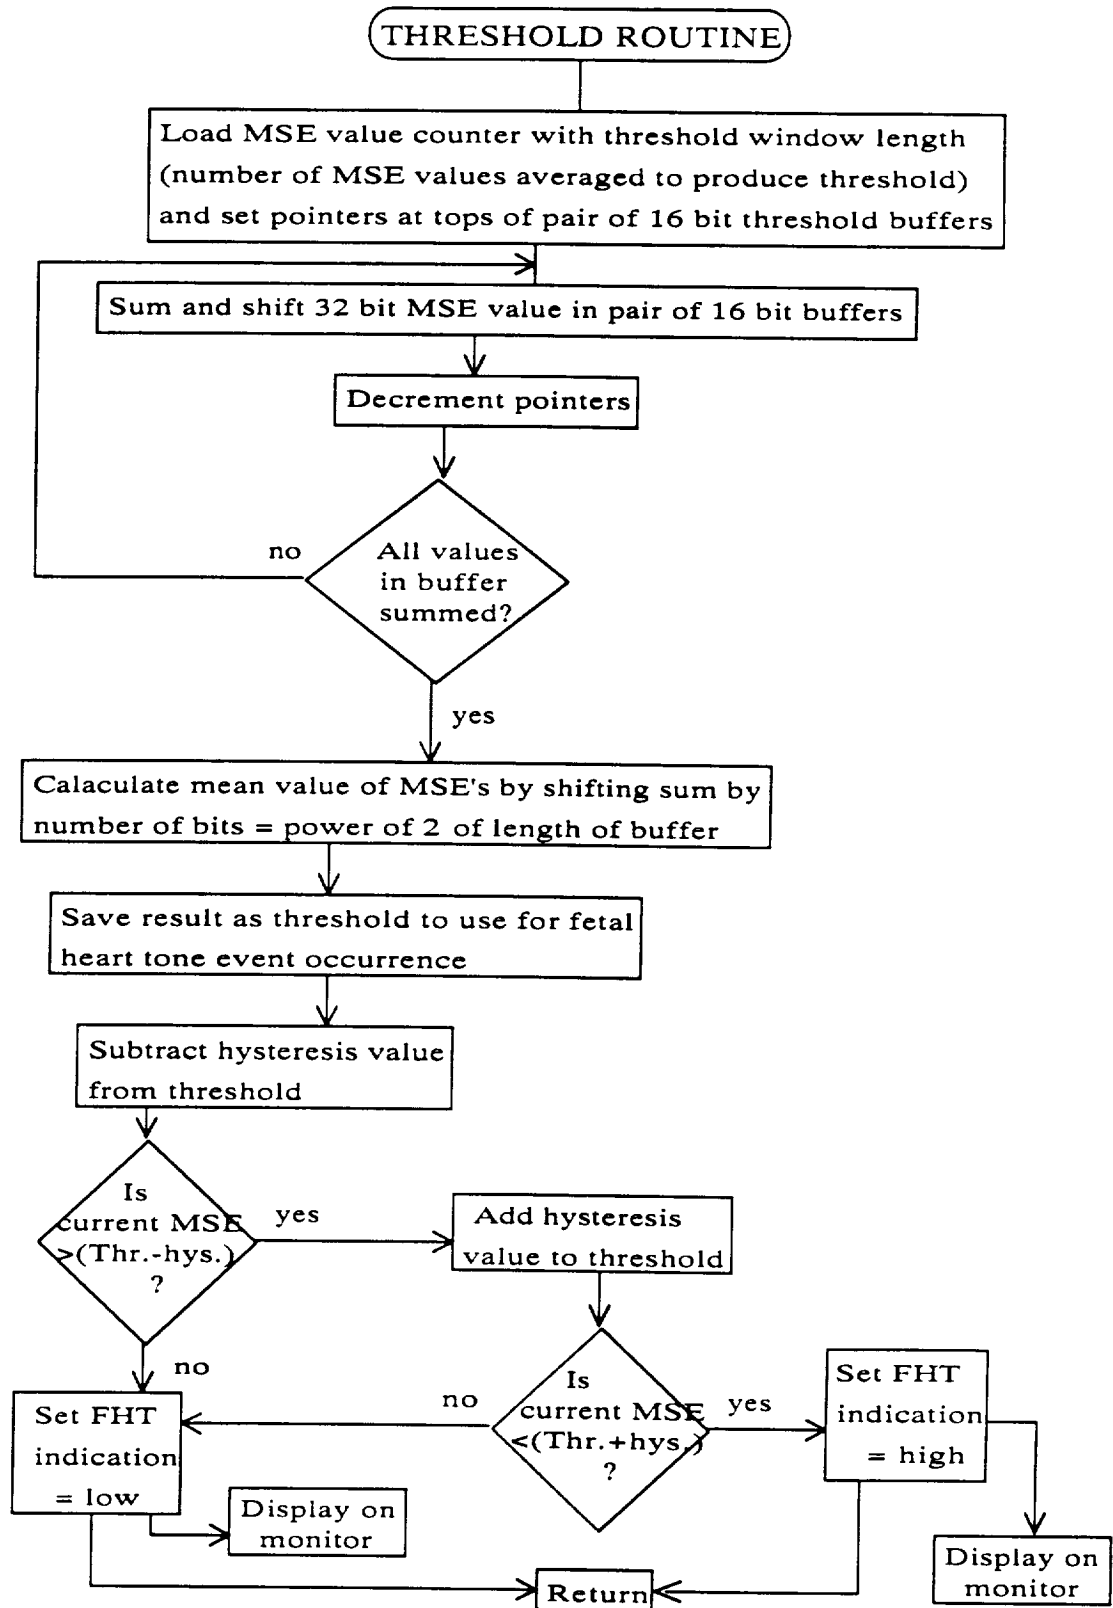
Figure 3.12. Threshold Routine.

information was difficult to determine, the hysteresis level in this research was set by empirical means.

If the current MSE is greater than the threshold, H_{MSE}, minus hysteresis *and* less than the threshold plus hysteresis, a FHT is judged to be present and the FHT indicator state is set high. Otherwise, the FHT indicator state is set low. To express this detection rule , if

$$(H_{MSE} - Hys)] < MSE < (H_{MSE} + Hys)$$

then FHT = High

Otherwise, FHT = Low

The FHT indicator state is plotted on the monitor display in real time below the MSE plot (figure 3.11). The heart rate calculation routine is then called.

3.4.7   Heart Rate Calculation Routine

The Heart Rate Routine is shown in figure 3.13. The fetal heart rate, R(n), in beats per minute is given by

$$R(n) = (60)(1 / T(n))$$

where T(n) is the period of the heart tones, i.e. the time interval between fiducial timing points, and n is the current sample number. The heart rate is updated with each new sample value. The fiducial timing points of the heart tones are the leading and trailing edges of the threshold-produced FHT indicator (high/low) states. A plot of the indicator

```
                          ( HEART RATE ROUTINE )

(Has a fetal heart
tone (FHT) edge
occurred?)                        ↓
                              ╱   is   ╲
                            ╱  previous  ╲
                          ╱ FHT state (high/low) ╲  yes    ┌─────────────────────────┐
                          ╲  same as new   ╱──────────────▶│ Increment leading and   │
                            ╲            ╱                  │ trailing edge interval  │
                              ╲  state? ╱                   │ counters; replace previous│
                                ╲    ╱                      │ FHT state with new state │
                                  ╲╱                        └─────────────────────────┘
(Therefore a FHT edge                                                  │
has occurred;  is it a      │ no                                       ↓
leading or trailing edge?)  │                                     ┌─────────┐
                            ↓                                     │ Return  │
                          ╱  is  ╲                                └─────────┘
                        ╱ new state ╲    no    ┌──────────────────────────────┐
                        ╲  a low?  ╱───────────▶│ Call leading edge procedure to│
                          ╲     ╱               │ calculate leading edge interval│
                            ╲ ╱                 └──────────────────────────────┘
                           │ yes                              │
                           ↓                                  │
              ┌─────────────────────────┐                     │
              │ Call trailing edge procedure to│              │
              │ calculate trailing edge interval│             │
              └─────────────────────────┘                     │
                           │                                  │
                           └──────────────┬───────────────────┘
                                          ↓
              ┌─────────────────────────────────┐     ╭──────────────────╮
              │ Average leading and trailing edge intervals│   │ Rate =  (f_s)(60) │
              │ and  calculate new instantaneous heart rate│──│         ─────────  │
              └─────────────────────────────────┘     │          count     │
                           ↓                           ╰──────────────────╯
              ┌─────────────────────────────────┐
              │ Call heart rate linear prediction procedure │
              └─────────────────────────────────┘
                           ↓
              ┌─────────────────────────────────┐
              │ Put final new instantaneous heart rate in buffer│
              │ and average with previous rates in buffer │
              └─────────────────────────────────┘
                           ↓
              ┌─────────────────────────────────┐
              │ Replace old FHT state with current state │
              └─────────────────────────────────┘
                           ↓
              ┌─────────────────────────────────┐       ┌──────────────┐
              │ Output new rate to display and strip │──▶│ Display on   │
              │ chart recorder (via D/A converter)│      │ monitor      │
              └─────────────────────────────────┘       └──────────────┘
                           ↓
                      ┌─────────┐            ┌──────────────────────┐
                      │ Return  │──────────▶│ Strip chart recorder │
                      └─────────┘            └──────────────────────┘
```
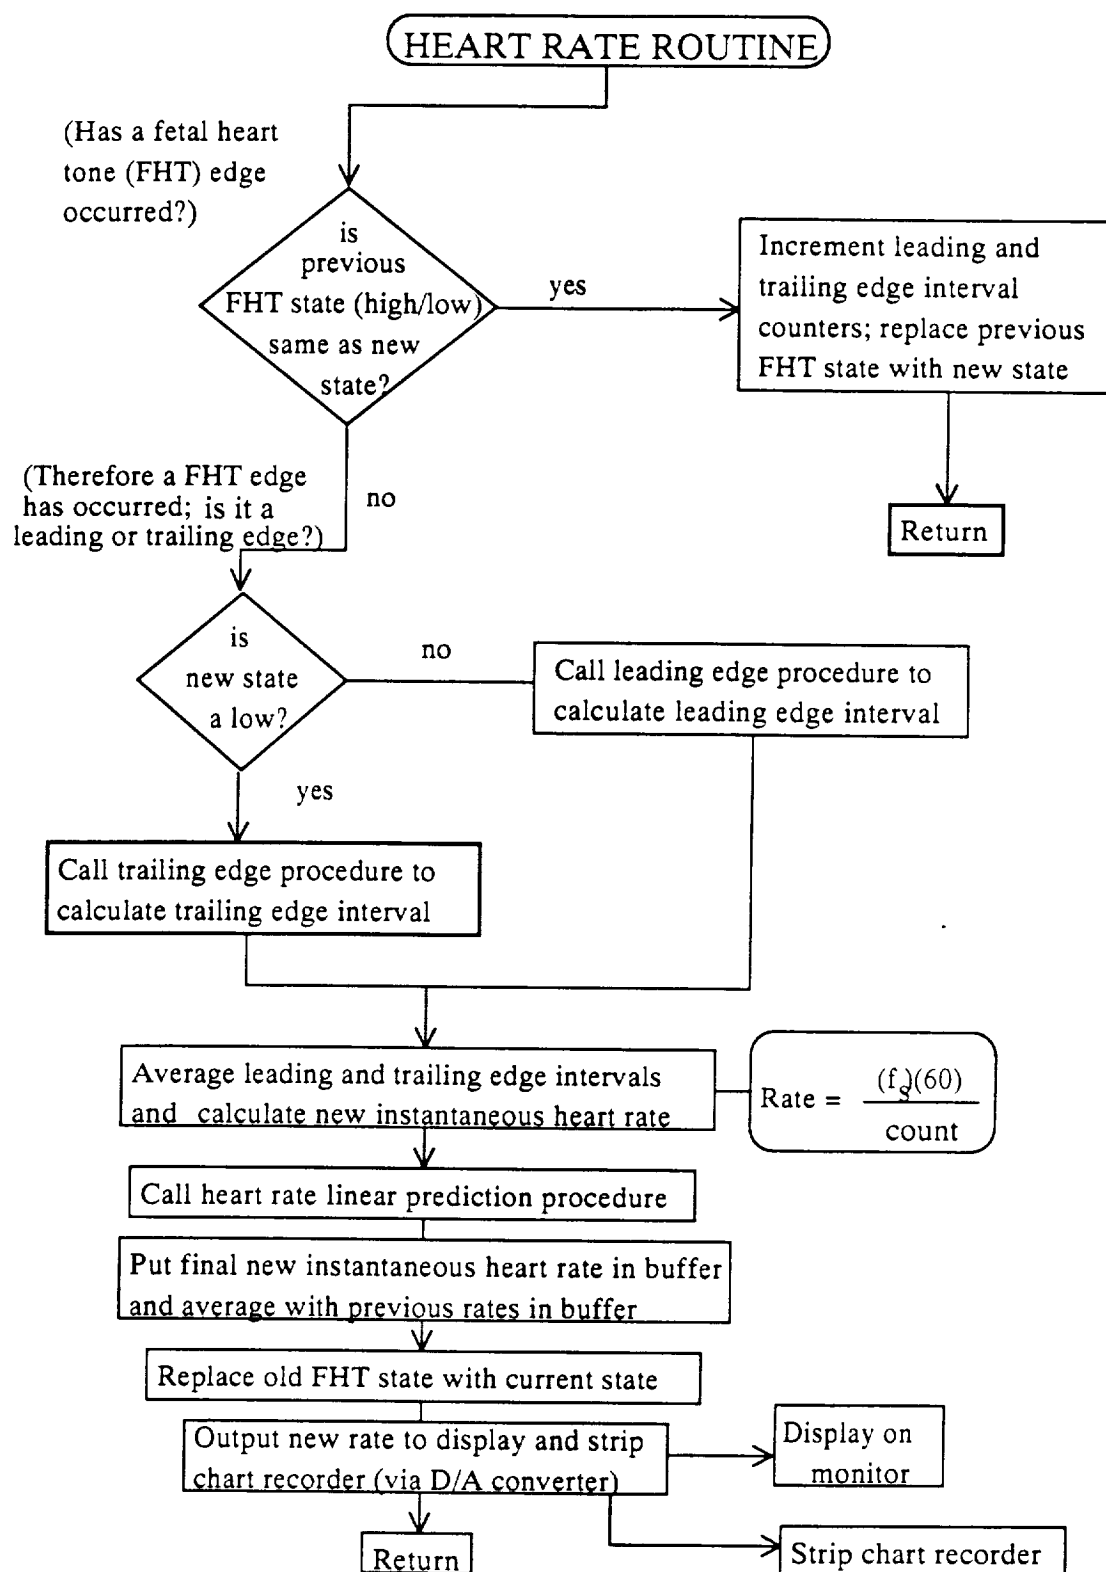
Figure 3.13. Heart Rate Routine

states has the appearance of a square wave, the period of which is the period of the heart tones (figure 3.11). To determine the period either the leading edge to leading edge interval or the trailing edge to trailing edge interval can be measured. The time of the interval, T(n), is calculated by

$$T(n) = S(n) / f_S$$

where S(n) is the number of samples in the interval and f_S is the sampling rate. The heart rate, R(n),in beats per minute,is therefore given by:

$$R(n) = [(f_S)(60)] / S(n)$$

In this algorithm both leading edge and trailing edge intervals are measured (by separate sample counters) in order to reduce error. An average of the two intervals is taken as the heart tone period.

The Heart Rate Routine must first determine whether a FHT edge has occurred. The current FHT state is compared with the immediately previous state. If the current state differs from the previous state, an edge has occurred. If they are the same an edge has not occurred and the edge to edge interval sample counters are incremented by one. The routine then returns to the Main Loop Routine for processing the next sample. If a FHT edge has occurred, the routine must determine whether the edge is a leading or trailing edge. If the new FHT state is a low, because it is different than the preceding state which must thus be a high, the edge is a trailing edge. Similarly, if the new state is a high, the edge must be a leading edge. A specific routine for each type of edge is then called. The edge routines are identical except for the variables and therefore will be discussed as one routine.

Leading and Trailing Edge Routines

A generic Edge Routine is depicted in figure 3.14. The Leading Edge and Trailing Edge routines measure the intervals between FHT leading edges and between FHT trailing edges, respectively by counting sample points . The routines are, however, also concerned with identification and correction of spurious intervals originating from missed heart tones or noise-induced false heart tones.

The interval sample count, $S(n)$, is compared to an expected range in order to identify spurious intervals. This range is based on the maximum rate of change for heart tone intervals observed in the data collection phase. Short term interval variability occurs due to blood flow dynamics. Long term variability occurs due to changes in overall heart rate, such as that accompanying fetal movement. The maximum rate of change from one interval to the next was observed (in the data collection pilot study) to be less than 20%.

The routine, therefore, uses an expected range of plus or minus 20% of the average of the immediately preceding edge to edge intervals, $A(n)$. If the current interval, $S(n)$, is within that range, the edge routine returns to the heart rate routine for calculation of heart rate as above. If the interval exceeds the range, the routine then substitutes an appropriate interval slanted toward the actual interval observed. To summarize, if

$$[A(n) - 0.2A(n)] \ < \ S(n) \ < \ [A(n) + 0.2A(n)]$$

then $S(n)$ is accepted.

Otherwise,

EDGE ROUTINE (generic)

(would have been set if last interval was spurious)

is rate-skip flag set?

yes → Reset rate-skip flag

no

Get edge interval count

is count outside of expected range?

no → Put "0" in signal quality buffer

yes

Substitute interval value = average interval slanted toward observed interval = av. int +-0.2(av. int.- observed int.)

Put "1" in signal quality buffer

Set rate-skip flag

Store count for heart rate calculation and put in count in buffer for range calculation

Increment other edge counter

Reinitialize same edge counter

Average previous edge counts in buffer

Update expected range for next count = average count +- 0.2(average count)

Reinitialize same edge interval count

Replace old FHT state with current state

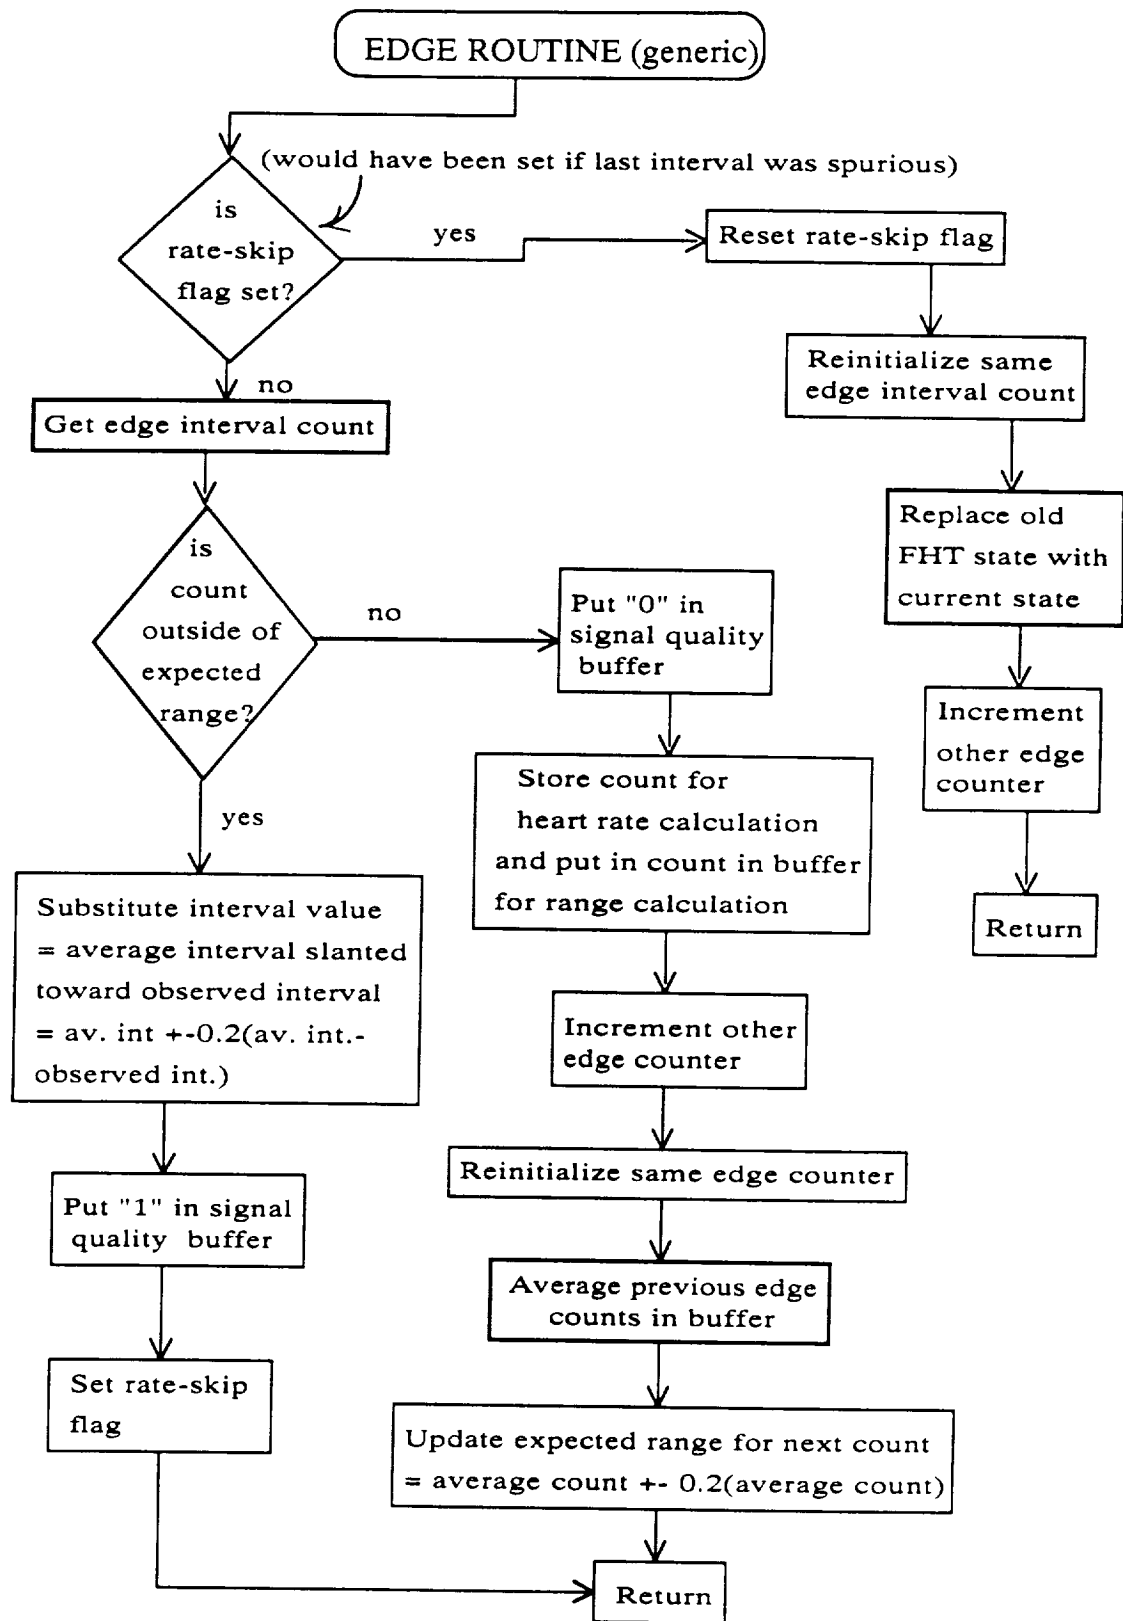Increment other edge counter

Return

Return

Figure 3.14. Edge Routine, Generic (Leading or Trailing).

$$S(n) = A(n) + C [S(n) - A(n)]$$

where C is the slanting factor. Slanting the interval toward the observed out of range value avoids "locking up" the substitution process. If an extremely small interval were encountered from a noise spike, for example, the average interval value could be pulled so low that all subsequent intervals would be greater than the range so that substitution would continue indefinitely. The sign resulting from the subtraction slants the substitution in the appropriate direction. A slanting factor of 0.2 was empirically used in this research.

Whenever a substitution is made, a "1" is put into a signal quality buffer location. This buffer is used by the Signal Quality Routine (discussed below) to keep track of the ratio of substituted versus good heart tone intervals as a measure of signal quality. Too many substitutions indicates poor signal quality and results in lifting of the strip chart pen. If the observed interval does not require substitution, a "0" is put in the signal quality buffer.

Because substitution is precipitated by encountering a spurious interval, there is no valid reference point for the start of the next interval. The edge routines therefore wait until the next edge of the same type is encountered before restarting the interval sample counters. This is done by setting an interval skip flag which holds counting until the next edge of the same type arrives. The remainder of each edge routine is involved with placing either the current or the substituted interval value in the previous interval buffer, taking the average of the buffer, and generating the new interval range value for comparison with the next encountered interval. The edge routines then return to the heart rate routine where the most recent leading edge and trailing edge intervals are averaged and the instantaneous heart rate is computed. This procedure is summarized by

$$R(n) = \frac{2(f_s)(60)}{S_L(n) + S_T(n)}.$$

where $R(n)$ is the instantaneous heart rate based on the current intervals and $S_L(n)$ and $S_T(n)$ are the leading and trailing edge interval sample counts, respectively.

Heart Rate Linear Predictor Routine

Figure 3.15 illustrates the Heart Rate Linear Predictor Routine. The instantaneous heart rate is computed from the two edge interval counts as discussed above. Spuriousness of one or the other interval counts requires a certain amount of time to recover via the slanting process. This can cause perturbations in the derived instantaneous heart rate. In an effort to minimize such perturbations a second linear predictor routine was used in the heart rate computation process. The heart rate linear predictor is described by

$$\hat{R}(n) = \sum_{k=1}^{M} R(n-k)P(k).$$

where $\hat{R}(n)$ is the predicted rate, $P(k)$ are the heart rate predictor coefficients, and M is the length of the predictor. Precalculated first order predictor coefficients were used which in effect fits the heart rate points to a straight line (including slope) in order to predict where the next rate point should be [25]. A pilot study evaluated various heart rate predictor lengths. A length of $M = 4$ was found to be most efficacious and was used in this research. Longer lengths tended to put the predictor into infinite substitution. The first order coefficients are:
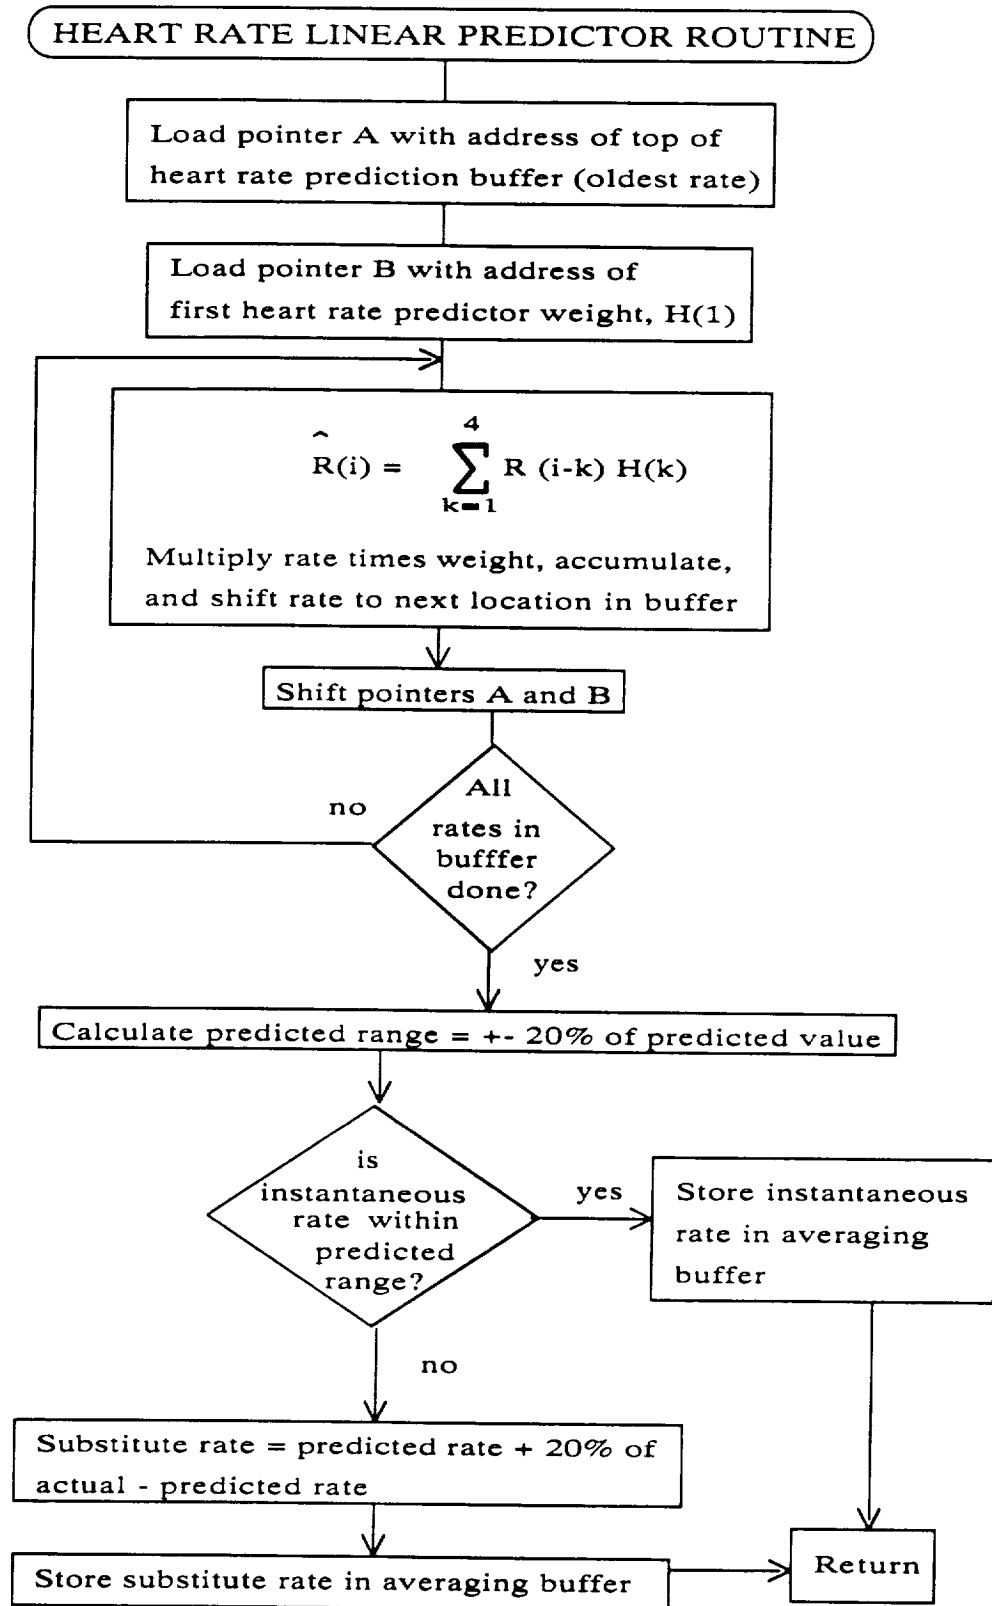
$$P(1) = -0.50$$

Figure 3.15. Heart Rate Linear Predictor Routine.

$$P(2) = 0.00$$
$$P(3) = 0.50$$
$$P(4) = 1.00$$

The current instantaneous heart rate, R(n), is compared to an expected range based on the predicted value. If the rate is within the expected range the instantaneous rate is stored for output to the strip chart recorder. Otherwise a rate substitution is made in a manner similar to that of the edge routine substitution process. The heart rate substitution procedure is as follows. If,

$$[\hat{R}(n) - G\hat{R}(n)] < R(n) < [\hat{R}(n) + G\hat{R}(n)]$$

then the current instantaneous heart rate is accepted.
Otherwise,

$$R(n) = \hat{R}(n) + J[R(n) - \hat{R}(n)]$$

G is a range factor, set at 0.2 based on the data collection pilot study analysis. J is a slanting factor set empirically at 0.2 for this research. The substituted rate is scaled for output via D/A converter #1 to the strip chart recorder.

The remaining portion of the heart rate routine is involved in further smoothing of the heart rate plot for output to the strip chart recorder. A keyboard selectable (optional) smoothing procedure is performed of the current and previous instantaneous rates over a block length B, where

$$R_{sm}(n) = 1/B \sum_{k=1}^{B} R(n).$$

For this research B is 16 rates.

3.4.8 Signal Quality Routine

The Signal Quality Routine, shown in figure 3.16, assesses the quality of the fetal heart signal being processed. The pen on the strip chart recorder is lifted if the quality becomes too poor. The edge interval routines described above insert either a "1" or a "0" into a rotating buffer contingent upon a substituted versus a valid heart tone interval. The Signal Quality Routine takes a sum of this buffer and compares that sum to a maximum number of substitutions allowed. If the proportion of substitutions is greater than the allowable maximum, a flag is set which when dectected by the PC brings about lifting of the strip chart pen and blocking out of the heart rate value on the monitor display with red. When the proportion of substitutions falls below the maximum again, the flag is reset. The PC then puts the pen back down and displays the heart rate again.

3.4.9 Sensor Select Routine

The sensor select routine (not illustrated) allows one of seven sensors to be selected by the TMS for input. A manual keystroke on the PC is detected by the TMS whereupon a pulse is sent out via D/A converter #2 to a counter chip controling the address input to the analog mutiplexer. The counter increments by one, selecting the next multiplexer channel. A routine for allowing the TMS to automatically select the optimal sensor was investigated. The minimum longterm mean MSE value appeared to be the best criteria for sensor selection. The automatic routine was not implemented in the real time system because of difficulties with fabricating a multisource acoustical generator required to test the software with the sensor belt. Future development of the system will include this feature.
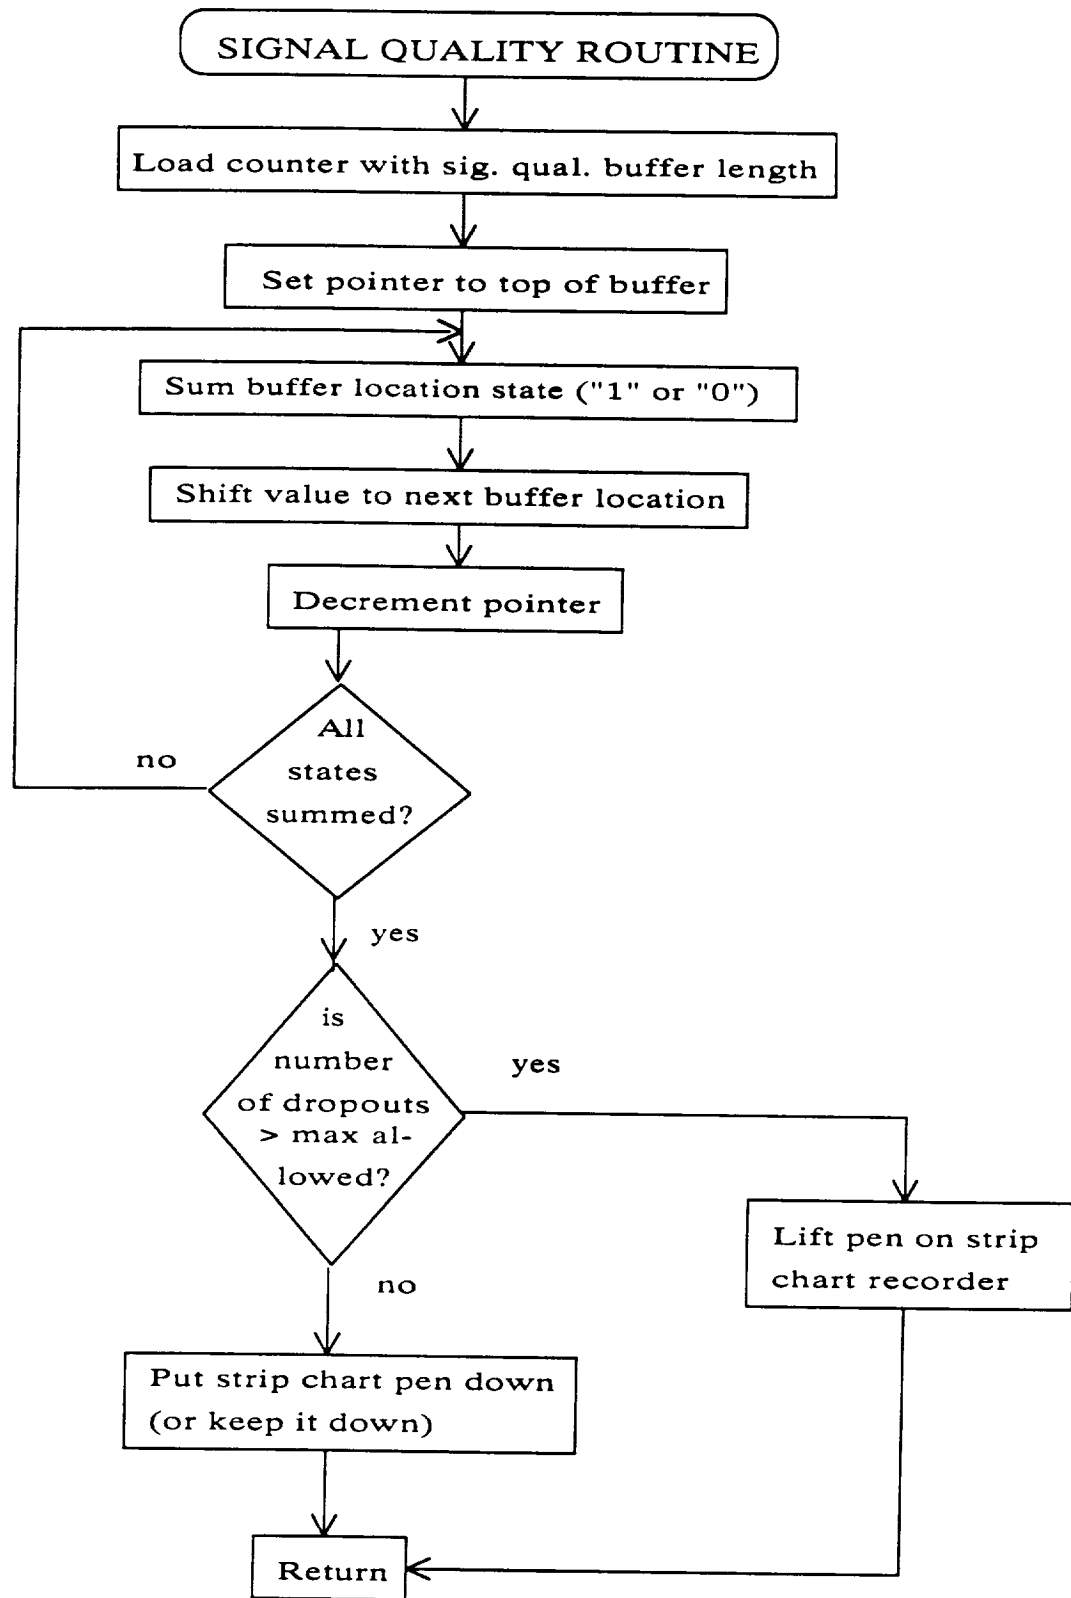
Figure 3.16. Signal Quality Routine.

3.4.10 Strip Chart Calibration Routine

The Strip Chart Calibration Routine (not illustrated) causes the TMS to output a full scale constant value to the strip chart recorder for calibration purposes upon detecting a certain keystroke on the PC keyboard . Once the key is released the TMS returns to the monitor algorithm.

3.5 Algorithm Variables

The real time fetal heart rate monitor algorithm of this research contains 15 variables listed in Table 3.1. Each of these variables was empirically adjusted using pilot studies. Optimization of these variables is planned for future development.

3.6 PC/TMS Handshaking Control/Display Routine

A handshaking program was devised to allow real time PC keyboard control of the TMS algorithm and to allow real time display of the incoming fetal heart signal and TMS generated parameters. The fetal signal and parameters such as MSE, FHT state, sensor number, overflow status, and heart rate are displayed in real time on the monitor screen as shown in figure 3.11. The PC/TMS Handshaking Control/Display Routine, shown in figure 3.17, halts the TMS, accesses its memory, reads and writes parameters, and then lets the TMS go while the output parameters are plotted on the monitor screen. The PC also controls the up/down position of the pen of the strip chart recorder depending on the signal quality flag of the TMS.

3.7 Processor Utilization

The processing time required for theTMS algorithm loop must not exceed the sampling period of the processor or data loss will occur. Confirmation that the total

Table 3.1. Fetal Heart Rate Monitor System Variables

| VARIABLE	CURRENT VALUE
1. Sampling rate	158
2. Predictor length	20
3. Error window for calculation of MSE	20
4. MSE window for calculation of threshold	128
5. Threshold hysteresis	7
6. Spurious FHT edge to edge interval substitution variables	
6.1 Number of previous edge-edge intervals averaged for next interval evaluation	4
6.2 Percent deviation of interval allowed before substitution made	20
6.3 Percent slanting of substituted value toward actual value	20
7. Heart rate linear predictor length	4
8. Heart rate substitution variables	
8.1 Percent deviation of heart rate allowed before substitution	20
8.2 Percent slanting of substituted value toward actual value	20
9. Pen lift variables	
9.1 Number of edge-edge intervals checked for substitutions	20
9.2 Ratio of good vs. substituted intervals which when exceeded initiates pen lift	10/20
10. Heart rate smoothing: number of instantaneous rates averaged	16
11. Anti-aliasing filter cutoff frequency (Hz)	55
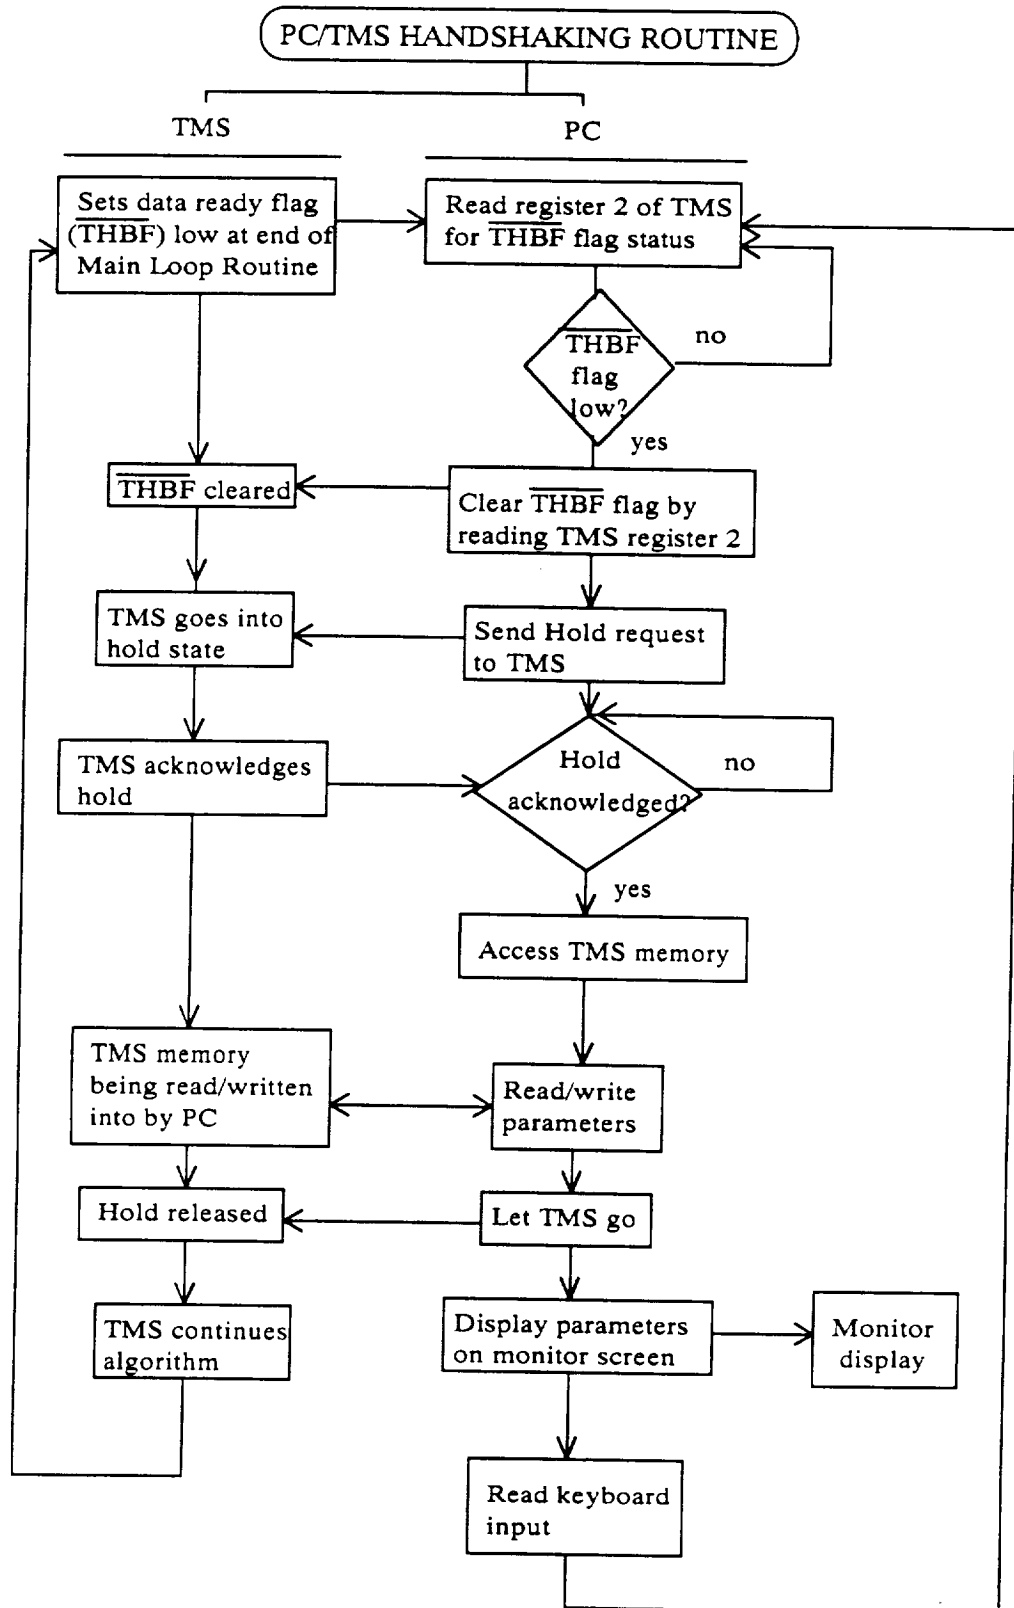
Figure 3.17. PC/TMS Handshaking Routine.

algorithm processing time was less than the sample period was accomplished by outputting to a D/A converter a high pulse transition at the start of the algorithm followed by a low transition at the end of the algorithm. The time between transitions, measured via an oscilloscope, gives the total algorithm time. It was required that the TMS board A/D converter be software driven for this procedure rather than letting the sampling rate clock control the converter. This is because the A/D and D/A are tied together so that only one D/A transition was possible if the sampling rate clock were allowed to control the A/D-D/A system. The TMS algorithm was found to require approximately 0.5 msec. for execution. With a sampling period of approximately 6 msec. this represents a processor utilization of 8 percent if the PC/TMS handshaking routine time is not included. The handshaking routine adds an additional 0.6 msec. which constitutes the time that the TMS is halted. This brings the total algorithm time to 1.1 msec. or 18.3 percent total utilization.

CHAPTER FOUR

EXPERIMENTAL STUDY

4.1 Introduction

The goal of this research was to develop a real time signal processing algorithm which would detect the time history of fetal heart tones within a noise contaminated acoustic signal, and from this information would derive heart rate. The real time methodologies to achieve this goal were presented in Chapter 3. The experimental study to validate the feasibilty of the LMS linear predictor algorithm and the associated heart rate derivation routine is described in the current chapter. The acoustic system is compared with a Corometrics Model 145 ultrasonic monitor in the clinical setting of the fetal NST.

The third generation sensor belt was developed for the comparative phase of this research, containing seven back-to-back sensor pairs in a 12 centimeter hexagonal array. It was hoped that the large arrray would allow tracking of the fetal heart tone point as it migrated with fetal movement.

4.2 Patient Subjects

Sixteen patients coming to the EVMS Department of Fetal Maternal Medicine participated in the experimental study. These were high risk mothers who were being monitored for fetal well being by weekly ultrasound NST's. The Institutional Review Board at EVMS gave approval for the second clinical phase of this research.

4.3 Procedure

On arrival at the EVMS clinic patients were asked by the NST technician whether they would be willing to participate in the acoustic monitor study. If they volunteered, written informed consent was obtained. The location of the fetal heart was then determined by means of an ultrasonic imaging device as per routine for the ultrasonic NST. A deLee Hillis fetal stethoscope was next used to locate the loudest heart tone point in the area of the fetal heart indicated by the imaging device. The acoustic and ultrasonic sensors were then positioned with the acoustic sensor being placed as close to the fetal heart tone point as the ultrasonic sensor position would allow.

Once the sensors were positioned the acoustic monitor system was powered up with concomitant initiation of the software. The strip chart recorder was calibrated by means of the PC keyboard and the TMS calibration routine (as described in section 3.3.10). The electronics were then powered up and the amplifier gain adjusted as necessary using the monitor screen to assess signal amplitude. The appearance of the real time monitor screen with an adequate fetal signal is shown in figure 4.1. Repositioning of the sensor was sometimes necessary to obtain an adequate signal. A software generated "beep" was output by the system as an audible indication of detection of each fetal heart tone. Once an adequate signal was confirmed, the strip chart recorder was initiated (figure 3.2). A mark was made on the ultrasonic strip chart to indicate where simultaneous acoustic monitoring began. The simultaneous recordings were carried out over the remainder of the NST, a period of approximately 15-20 minutes after which both sensors were removed.

The acoustic and ultrasonic sensors competed for the same location on the mother's abdomen. For obvious medical reasons the ultrasonic sensor had priority. Because of this site competition it was not possible to obtain an adequate fetal acoustic signal from six patients in the study and in no patient was it possible to use the third generation array sensor belt due to its greater width. All of the simultaneous comparative

recordings were therefore done with the second generation belt. Figure 4.2 shows the second generation sensor belt and the ultrasonic sensor in place along with an additional sensor used for uterine contraction monitoring. The third generation belt was placed for additional heart rate recording after the ultrasonic NST if the patient was agreeable and the fetus was stable. Indirect confirmation that the third generation belt would produce the same results as the second generation belt was thereby possible. The third generation belt is shown in figure 4.3 and is shown implemented in figure 4.4. The mother is holding a button that is pushed whenever fetal movement is felt.

One other problem resulted from the close proximity of the acoustic and ultrasonic sensors. It is common practice for the technician performing the NST to manually hold the ultrasonic sensor in position. Vibrations from the technician's hand tended to be picked up by the acoustic sensor causing untenable signal to noise ratios in four of the patients in the study. The comparative results from the remaining six patients are presented in the next section.

4.4 Ultrasound Comparative Study

The comparative acoustic and ultrasonic heart rate strip chart records for six patients are shown in figures 4-5 through 4-10. The acoustic heart rate record (shown in red) was superimposed via a digital scanner onto the ultrasonic record (shown in black) for direct visual comparison. Because the ultrasonic numerical data was unavailable, only anecdotal comparison is possible.

Fetal movements are indicated by vertical black arrows with "FM" below them. Note that the best correlation between the acoustic and ultrasonic records occurs during heart rate accelerations associated with fetal movements.

Figure 4.11 is a heart rate record produced with the third generation multi-sensor belt alone. This record was immediately subsequent to a comparative record with the second generation belt system versus the ultrasonic unit shown in figure 4.10. The third

Figure 4.4 Third Generation Sensor Belt Implementation

PATIENT 23

ULTRASONIC   ACOUSTICAL

PATIENT 22B

ULTRASONIC ——— ACOUSTICAL

Figure 4.6. Acoustical vs. Ultrasonic NST : Patient 22b.

PATIENT 16

ULTRASONIC ———— ACOUSTICAL

Figure 4.7 Acoustical vs. Ultrasonic NST : Patient 16.

PATIENT 17

ULTRASONIC ——————   ACOUSTICAL

Figure 4.8.   Acoustical vs. Ultrasonic NST : Patient 17.

PATIENT 30

ULTRASONIC ———— ACOUSTICAL

Figure 4.9. Acoustical vs. Ultrasonic NST : Patient 30.

PATIENT 31

ULTRASONIC ——— ACOUSTICAL

Figure 4.10.   Acoustical vs. Ultrasonic NST : Patient 31.

generation heart rate record shows a clear acceleration associated with fetal movements denoted by the small spike on the tracing produced by the mother's pushing the hand button. It is concluded that the third generation belt would perform comparably to the second generation belt if a comparison test with the ultrasound unit were possible.

4.5 Discussion of Results

The purpose of this experimental study was to test the feasibilty of the acoustic system and the linear predictor algorithm for real time fetal heart tone detection and heart rate derivation by comparison with an industry standard ultrasonic monitor. There is good anecdotal correlation between the ultrasonic and acoustic heart rate records in five of the six patients whose data is presented. The results from the other ten patients in the study do not constitute valid comparisons due to disruptive extraneous factors. The correlation is strongest during the periods of heart rate accelerations, a result possibly attributable to the fact that fetal heart tones are of greater amplitude during accelerations. Drop out (pen lift) periods are attributable to factors such as maternal movement, changing fetal position, or excessive background noise.

CHAPTER FIVE

CONCLUSION

5.1 Overview

This research has demonstrated the feasibilty of the LMS linear predictor algorithm for the real time detection of fetal heart tones from a noise contaminated acoustic signal with subsequent derivation of heart rate. The acoustic linear predictor system compares favorably with a commercial ultrasonic monitor for heart rate generation in the context of the fetal nonstress test. Certain problems remain to be solved, as discussed below. The system also contains multiple variables which require optimization.

5.2 Algorithm Performance

1) Good anecdotal correlation in five out of the six valid comparative studies indicates a high degree of efficacy for detection of fetal heart tones by the linear predictor algorithm.

2) The algorithm is highly time efficient requiring only 18% processor utilization for a TMS320C25 running at 40 MHz.

3) The algorithm is space efficient requiring less than four kilobytes of TMS board memory.

4) Problems encountered:

   The most significant problem encountered with the linear predictor algorithm was its sensitivity to noise, particularly maternal movement artifact. It is

conceivable that the fetal heart tone sound can be mimicked by bodily "thumps" such as the mother moving her hand against her leg. But gross observation of the predictor performance does not indicate such. One reason for the predictor's noise sensitivity may be that there is movement artifact included in the training data from which the predictor weights were derived. In the time domain whole file training process used for this research the updating of the weights is a function only of local signal power and therefore would also train on noise spikes.

Other problems encountered include sensitivity of the predictor to maternal heart tones. This problem was significant in the mothers with large overlap in the fetal-maternal spectrum.

5.3 Other System Performance Aspects

1) The electronic components and parameters such as amplifier gain and filter cutoffs appear to be adequate.

2) The sensor belt suffers from two problems:

    a) The sensors are marvelous accelerometers and therefore quite sensitive to rigid body motion. Compare the sensors to a simple bell stethoscope which allows detection of the fetal heart tones with little or no body motion sensitivity.

    b) The sensor sensitivity to the fetal signal is somewhat low resulting in a low signal to noise ratio. A gain factor of at least 5,000 is required to amplify the fetal signal to the 3 volt amplitude level required to make use of the dynamic range of the TMS board A/D converter. Compare this again to a stethoscope which requires no amplification.

5.4 Future algorithm/system development potential

1) Optimization of the fifteen algorithm variables listed in Table 3.1 may significantly enhance the performance of the real time system.

2) Alternate predictor training methods may make the predictor more signature specific and less noise sensitive. One training improvement would be to add a detector to the time domain whole file training algorithm. The weights would be updated only when both the local energy level exceeds a threshold value and a heart tone is also detected by the predictor using a generic set of weights. Still another training improvement would be real time training on each new patient. Processing of the patient signal would be initiated with a generic set of weights and these initial weights would be continually updated during processing. Again, an effective method is needed to turn on the weight updating only when a heart tone is encountered and off between tones.

3) More precise cancellation of accelerometer-type motion artifact could significantly enhance the signal to noise ratio at the front end of the system. Differential addition of the out of phase signals from the back to back sensor pair still appears to be the best approach to this goal. But in order for this concept to work properly the exact proportions of the acceleration noise signal on each of the two sensors of the pair has to be known. Then the sensor outputs can be appropriately scaled before differential addition. The ratio of acceleration sensitivity of the pair may, however, not be a constant and may be a function of such factors as belt tightness. Adaptive noise calcellation techniques could then be used.

4) Better characterization of the fetal heart tone signal may allow more specific detection methods to be devised. For example, instead of being a single system response, does the fetal heart acoustic signal represent the response of several sub-systems as manifested by the first and second heart sounds or the M,T,A,P sub-components. These subsytem responses may not be interdependent. A predictor trained on the composite system response may fail when phase differences between the sub-system responses occur. Perhaps a solution is to use multiple predictors, each trained on a different sub-system response (ex. the first or second heart sound or a single sub-component). These multiple predictors would operate concurrently with a requirement for detection being that all the predictors simultaneously reach MSE minimums.

5) A possible method to increase the signal ouput of the sensors is to generate tension on the sensor surface. The second generation belt has twice the signal output of the third gereration belt, possibly because its sensors became bent from repeated tight contact with the maternal abdomen.

5.6 Future Development Potential of the Acoustic Monitor Concept

The acoustic fetal heart monitor concept has several potential development areas:

1) Acoustic detection could be extended to detection of fetal movement in a manner similar to that of detection of the heart tone signature. If fetal movement could be detected in addition to the heart tones, then automation of the NST would be possible. The acoustic sensors have a fairly flat frequency response down to DC and could be used to detect even slow fetal movements. The monitor could conceivably detect fetal movements which are imperceptable to the mother.

2) Automatic sensor scanning using the third generation multiple sensor belt is a planned development. The sensor with the minimum long term MSE level would be the theoretical optimum sensor.

3) Minaturization of the system could allow home use of the acoustic monitor. This would enable greater surveillance of high risk pregnancies and provide convenience for the mother.

The totally non-invasive nature of the acoustic system lends itself well to long term monitoring of high risk pregnancies. Detection and prevention of the causes of such devastating maladies as cerebral palsy and birth defects are ostensibly within the realm of possibilities. This research has demonstrated the feasibility of an idea. The acoustic linear predictor fetal heart rate monitor has been shown to work. Continued effort should bring the idea to a useful reality.

REFERENCES

1. Goodlin, R. C., " History of fetal monitoring," *Amer. J. Obstet. Gynecol.,* vol. 33, no. 3, pp. 323-352.

2. Kennedy, E., *Observations on Obstetric Auscultation ,* Hodges and Smith, Dublin, Ireland, 1833.

3. Parer, J. T., *Handbook of Fetal Monitoring ,* Saunders, 1983.

4. "Nonstress FHR Testing," *Clinical Obstetrics and Gynecology,* vol. 25, no. 4, pp. 689-705, Dec. 1982.

5. Solum, T., "A comparison of three methods for external fetal cardiography," *Acta Obstect. Gynecol. Scand.,* vol. 59, 1980, pp.123-126.

6. Talbert, D. G., Davies, W. L., Johnson, F., Abraham, N., Colley, N. Southall, D. P., "Wide bandwith fetal phonography using a sensor matched to the compliance of the mother's abdominal wall," *IEEE Transactions on Biomedical Engineering,* vol. BME-33, no. 2, February 1986.

7. Hewlett-Packard GMBH, Operating and Service Manual for Model 8020A Cardiotacograph, Hewlett-Packard GMBH, Boblingen, West Germany, 1970.

8. Jenssen, H., "Fetal systolic time intervals after paracervical block during labor," *Acta Obstet. Gynecol. Scand.,* vol. 59, pp. 115-121, 1980.

9. Talbert, D. G., Dewhurst, J., Southall, D. P., "New transducer for detecting fetal heart sounds: Use of compliance matching for maximum sound transfer," *Lancet,* February 25, 1984, pp. 426-427.

10. Kobayashi, K., Yasuda, T., and Suzuki, K., "Piezoelectric highpolymer film for fetal phonocardiograph and the monitor," *Journal of the Acoustical Society of America,* vol. 64, Supplement no. 1, Fall 1978.

11. Nagel, J., "New diagnostic and technical aspects of fetal phonocardiography," *European Journal of Obstectrics and Gynecology and Reproductive Biology,* vol. 23, December 1986, pp. 295-303.

12. Ferrara, E., Widrow, B., "Fetal electrocardiogram enhancement by time-sequenced adaptive filtering," *IEEE Transactions on Biomedical Engineering,* vol. BME-29, no. 6, June 1982, pp.458-460.

13. Durand, L., "Digital signal processing of the phonocardiogram," *IEEE Engineering in Medicine and Biology Society 9th Annual International Conference Proceedings,* 1987, pp. 906-907.

14. Cohen, A., *Biomedical Signal Processing,* vols. 1 & 2, CRC Press, Boca Raton, Florida, 1986.

15. Lee, C., Wei, L., " Spectrum analysis of human pulse," *IEEE Transactions on Biomedical Engineering,* vol. BME-30, no. 6, June 1983, pp.348-352.

16. Azevedo, S., Longini, R., "Abdominal-lead fetal electrocardiographic R-wave enhancement for heart rate determination," *IEEE Transactions on Biomedical Engineering,* vol. BME-27, no. 5, May 1980, pp.255-260.

17. Manning, G. K., Dripps, J. H., "Comparison of correlation and modulus difference processing algorithms for the determination of fetal heart from ultrasonic Doppler signals," *Medical & Biological Engineering and Computing,* vol. 24, March 1986, pp. 121-129.

18. Fukushuma, T., Flores, C., Hon, E., Davidson, E., "Limitations of autocorrelation in fetal heart rate monitoring," *American Journal of Obstetrics and Gynecology,* vol. 153, 1985, pp. 685-692.

19. Hamilton, P., and Thompkins, W., "Detection of ventricular fibrillation and tacycardia by adaptive modeling," *IEEE Engineering in Medicine and Biology Society 9th Annual International Conference Proceedings,* 1987, pp.1881-1882.

20. Lin, K., and Chang, W., "ECG signal analysis by linear predictive method," *IEEE Engineering in Medicine and Biology Society 9th Annual International Conference Proceedings,* 1987, pp.557-558.

21. Venkat, Subhashri, *Visual Speech Training Aid for the Deaf,* Masters Thesis, Old Dominion University, 1990.

22. Alexander, T., *Adaptive Signal Processing: Theory and Applications,* Springer-Verlag, New York, 1986.

23. Widrow, B., Stearns, S., *Adaptive Signal Processing,* Prentice Hall, Englewood Cliffs, New Jersey, 1985.

24. McLennan, C., E., *Synopsis of Obstetrics,* C. V. Mosby Company, St.Louis, 1970.

25. Garadia, D., H., *Deterministic Linear Predictors for Bit Rate Improvement*, Masters Thesis, Old Dominion University, 1986.